Handelshøyskolen BI

GRA 19703 Master Thesis

Thesis Master of Science 100% - W

Predefinert informasjon

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

Deltaker

Navn:	**Andreas Becker Cappelen og Christoffer Thalberg Hamnes**

Informasjon fra deltaker

Tittel *:	Effects of Norwegian Oil Supply Shocks on the Global Oil Market: Implications for Environmental Policies
Navn på veileder *:	Thomas Størdal Gundersen

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

Gruppe

Gruppenavn:	(Anonymisert)
Gruppenummer:	167
Andre medlemmer i gruppen:	

Christoffer Thalberg Hamnes
Andreas Becker Cappelen

BI Norwegian Business School – Thesis

Effects of Norwegian Oil Supply Shocks on the Global Oil Market: Implications for Environmental Policies

Delivery date:
09.01.2023

Submission date:
03.07.2023

Supervisor:
Thomas Størdal Gundersen

Master of Science in Business, major in economics

Acknowledgements

We want to express our gratitude to Thomas Størdal Gundersen, our supervisor, for his valuable feedback and guidance throughout the development of this thesis. We would also like to thank Maximilian Schröder for his motivation and insightful discussions regarding the conditional forecast approach. Also, we would like to extend our appreciation to our lecturers and professors for their continuous support and encouragement in pursuing this field of study, along with our classmates for their engaging discussions and valuable insights.

Oslo, 21th of June 2023

Andreas Cappelen Christoffer T. Hamnes

Abstract

We examine the effect of Norwegian supply shocks on the global oil market and evaluate the environmental consequences. In our research question, we seek to find out how the global oil market will react if we have a shutdown of Norwegian oil production. The background of Norwegian oil production is discussed, along with previous and ongoing debates on the topic. A shutdown can have different effects depending on whether it is temporary or permanent. To evaluate the effect of a temporary shutdown of Norwegian oil production, we first look at oil market variables after a strike on Norwegian oil platforms in 1986. The strike resulted in a price increase of 27%. Second, we generalize the mechanism by constructing a structural VAR model. Our results show that a shock similar in size to the 1986 strike in today's context raises the oil price by 6.5%. Long-term global oil production was unchanged in both cases. We apply a conditional forecast to analyse the effects of permanent shocks to Norwegian oil production. Our results suggest that the oil price responds greatly to this shock, while global oil production increases in the first nine months before decreasing and stabilizing at a lower level. We find that less oil is produced in the shutdown scenario, implying that a shutdown may lead to beneficial outcomes for reducing greenhouse gas emissions.

Contents

1 Introduction

Crude oil is a vital source of energy that powers the global economy. In recent years, there has been growing awareness of the urgent need to transition away from fossil fuels to mitigate the harmful effects of climate change. As an oil-producing country, Norway plays an important role in this transition.

Fossil fuels are the largest contributors to climate change, accounting for more than 75 percent of greenhouse gas emissions (United Nations, n.d). To achieve the objective of limiting global warming to below 2°C as set forth in the Paris Agreement, it is necessary to acknowledge that a significant proportion of the world's fossil fuel reserves cannot be burned and must remain unextracted. The problem is that the costs of taking action to mitigate climate change, such as transitioning to renewable energy sources or reducing own production, are borne by each country. At the same time, the benefits of a healthier environment are shared by all countries. Since there are no global regulations to lower fossil fuel extraction, countries have limited incentives to act. In this thesis, we are trying to answer if Norway nevertheless can have an effect on global oil production and thus global CO2 emissions by shutting down domestic production.

Our research question is: How will the global oil market respond if Norwegian oil production were to be shut down? We will use empirical results to argue whether Norwegian climate actions can create a positive impact on greenhouse gas emission reduction. We hypothesize that a Norwegian supply shock will encourage higher oil prices but leave no significant reactions from global producers, reducing net emissions to some extent.

To evaluate our hypothesis, this thesis includes an analysis of a strike on Norwegian oil platforms in 1986, a structural vector autoregressive (SVAR) model, and a conditional forecast (CF). The strike, caused by kitchen workers on Norway's offshore oil rigs, led to a shutdown of substantial parts of Norwegian production. This gives us the opportunity to analyze how oil variables react to an exogenous Norwegian supply shock, giving plausibility to the interpretation of our analyses. Our structural VAR is constructed to answer if we can expect the same reactions conditional on a data sample running from 1st of January 1974 to 31st of December 2019, as we did in 1986. By employing a conditional forecast, we can simulate a persistent event where Norwegian production is constrained to zero on impact. This allows us to investigate the effects on global production and the price of oil.

Our findings show that a permanent shutdown of Norwegian oil production will have no effects on CO_2 emissions in the short run but have a beneficial impact from around nine months after the shutdown. The real price of oil is highly affected by a shutdown and increases from 22 $/bbl to around 90 $/bbl over a two-year period, before decreasing again. We also find that a temporary shutdown of Norwegian oil production results in short-term increased oil prices and no long-term response in global production.

The following section covers Norwegian oil history, relevant literature and recent debates, before briefly discussing supply vs. demand, and lastly, temporary vs. permanent supply shocks. Section 3 is about the 1986 strike. Section 4 includes the method, approach, and results for both the structural VAR and the conditional forecast. Lastly, in Section 5, we conclude and motivate further research.

2 Background and important research

2.1 Background

Since production on the Norwegian continental shelf started in the early 1970s, the value of oil and gas production has contributed over 18,000 billion to Norway's gross domestic product measured in today's currency (Norsk Petroleum, 2023). Since 1996 the income from the Norwegian petroleum industry has been transferred to The Government Pension Fund, which invests the money (Norges Bank, 2019). The money provides a stable source of income for the government that is used to support public spending and investment projects, including infrastructure, education, healthcare, and public pensions. The fund aims to ensure that both current and future generations benefit from petroleum revenues. For 2023, the Norwegian government will have an estimated net cash flow of around NOK 1,384 billion from the oil industry (Norsk Petroleum, 2023). The peak of Norwegian oil production was in 2001. Since then, production has been approximately halved (see Appendix A.2, Figure 9). However, total petroleum production has remained relatively constant due to increased natural dry gas production. In 2021/2022, 94 Norwegian oil fields were in use, and 88 new fields were assessed for potential oil extraction. The petroleum sector is by far Norway's most significant industry, employing around 200,000 people. It is estimated that only half of the total recoverable resources on the Norwegian continental shelf have been extracted, meaning that Norway has the opportunity to continue its petroleum production for many more years.

These benefits come with an environmental cost. Burning fossil fuels releases carbon dioxide and other greenhouse gasses into the atmosphere. These gases trap heat and contribute to the warming of the Earth's surface. In 2020, crude oil alone contributed 32% of global greenhouse gas emissions (IEA, 2021). In light of these environmental impacts, there have been calls for a reduction in oil consumption. Ryggvik and Rosendal (2021) claimed that if Norway wants to be an active contributor to reaching the UN's climate goals, down-scaling of petroleum activity on the Norwegian continental shelf should start immediately. They suggest that Norway should put an end to new explorations of oil fields and that the government should tax oil companies harder. Among the political parties in parliament, there are divided views on petroleum production. Parties like Arbeiderpartiet and Høyre want to develop and adapt the petroleum industry, not discontinue it, while Fremsktittspartiet wants an expansion (NRK, 2021). Venstre, Sosialistisk Venstreparti, Rødt, and MDG are among those parties that want Norway to stop looking for new oil fields. They want to remove tax benefits for oil companies and force cuts in emissions from oil and gas production (NRK, 2021). Our thesis seeks to fit into this debate.

Although the environmental costs of fossil fuels are large, there are several reasons for Norway to continue oil and gas production. The first one is the obvious monetary aspect discussed previously. Huge parts of the Norwegian economy are built around oil, and income from export is essential for Norway's current welfare. Secondly, Norwegian petroleum contributes to geopolitical stability. The majority of the world's oil and gas resources are controlled by countries without democratic governance. By moving production from Norway to these countries, we give more power and influence to non-democratic countries. Further, the war between Russia and Ukraine in 2022, and the following energy crisis in Europe, have highlighted the importance of energy self-sufficiency. Another dimension is that the petroleum industry in Norway is important for the development of new technologies, especially in the energy sector. The Norwegian Ministry of Petroleum and Energy highly encourages research and development, and companies can get support and funds from The Research Council of Norway. Technological breakthroughs in the energy sector are extremely important for a green future, so it is essential that oil production is carried out in countries that reinvest in R&D. Lastly, there is the theoretical possibility that a stop in Norwegian petroleum production does not affect global CO2 emissions due to increased supply by other producers.

2.2 Related literature and ongoing debates

The standard approach to address the externality problem associated with fossil fuel use has typically been through taxation on demand. Such a tax is normally a fee on carbon emissions from the production of goods and services, often referred to as the Pigouvian tax. The intention is to highlight the hidden social costs of greenhouse gas emissions and, in turn, reduce demand for products that are carbon-intensive. Other demand-side policies are the EU's Emissions Trading System (EU ETS) and Carbon Border Adjustment Mechanism (CBAM) (European Counsil, 2022). There has previously been relatively little focus in the literature on supply-side policies for fossil fuels compared to demand-side policies. However, recently there has been an increase in theoretical research on limiting fossil fuel production (Ahlvik, 2022), and we want to contribute to this research.

Fæhn et al. (2017) have previously estimated that a halving of Norwegian oil production could lead to an oil price increase of around one percent in the long-term. The authors write that a higher oil price will reduce the demand through direct and indirect effects. For oil producers, a higher price will make oil extraction more profitable, and supply will increase. These two forces are working in opposite directions, and the outcome is a result of price elasticity in the oil market. The sensitivity of oil demand will depend on the price of close substitutes for oil consumption, such as coal, gas, and renewable energy, as well as people's opportunity to reduce energy use. For supply, the sensitivity depends on the cost of searching for and developing new profitable oil fields and the costs of increasing production on existing fields. The literature does not give an unambiguous answer to the different elasticities in the oil market, and the estimates vary quite a lot. In Fæhn et al.'s. study from 2017, they review a substantial part of the empirical literature on long-run elasticities. While emphasizing that a consensus estimate of the different long-run elasticities in the oil market is difficult to nail down, they landed on a benchmark demand elasticity of -0.5 and a supply elasticity for non-OPEC countries of 0,5 (Fæhn et al., 2017). There is also the issue that the results in these studies are highly sensitive to the choices of elasticities. These authors also consider that a reduction in oil production will be met by increased consumption of other types of energy and that the environmental result will depend on what type of energy this is. The article concludes that a total shutdown of Norway's oil production would increase oil prices by about 2 percent in the long-term and that the price effect would be slightly larger in the short-term. The paper also suggests that the most cost-effective approach for achieving global reductions in greenhouse gas emissions within a domestic context would be to shift approximately two-thirds of planned demand-side abatement efforts to supply-side measures, specifically reducing oil extraction (Fæhn et al., 2017).

Rystad Energy (2023) recently published a report challenging Fæhn's (2017) findings and concluding that increased Norwegian oil production results in lower global greenhouse gas emissions per barrel. Rystad's calculated effects from increased future oil production on the Norwegian continental shelf show that approximately 90% of production goes to replace or out-compete other oil production, while 10% is absorbed through increased demand. In the report, the emission effect of increased Norwegian oil production is broken down into the following three steps:

1) Increased demand (10%) will contribute to increased emissions from end-use and combustion of oil, equivalent to 42 kg of CO2 per barrel of new oil production delivered to the market.[1]
2) Increased oil demand (10%) is expected to displace electric vehicle use due to lower oil prices. The avoided emissions from reduced power generation for transport amounts to 16 kg of CO2 per barrel of new oil production delivered to the market.
3) Emissions from increased Norwegian oil production (with electricity from shore) and decreased foreign outcompeted oil production yield a reduction of 52 kg of CO2e[2] per barrel.

Taken together, steps 1–3 results in a net global emissions reduction of 26 kg of CO2e per barrel. The achieved climate effect is a combination of low demand response in the oil market and significantly lower emissions from Norwegian production compared to the outcompeted oil production (Rystad, 2023). The report concludes that increased Norwegian gas production will partially replace coal consumption, resulting in even lower total greenhouse gas emissions.

The credibility of Rystad's report has been questioned on the basis that it was commissioned by the Norwegian Ministry of Petroleum and Energy. As a counter, the environmental protection organization WWF, Naturvernforbundet, Natur og Ungdom, and Greenpeace came together and commissioned Vista Analyse AS to write their own report on the same matter. Vista's report was published in March 2023, one month after Rystad's, and concluded that increased Norwegian oil production results in a net increase in global emissions, regardless of scenario and time horizon (Riekeles, 2023). Their findings show that the net increase in emissions is 47 kg CO2e per barrel of oil extracted in Norway in the baseline model and 90 kg CO2e in the low-emission scenario. This shows

[1] By comparison, burning one barrel of oil emits 419 kg of CO2.
[2] CO2e is a collective term for all greenhouse gasses.

that calculating the net emission effects of increased Norwegian oil and gas production is complex. The effects depend on several uncertain factors, and different estimates of these will give different results. Such as differences in the computation of consumption effects, the influence on other producers, the types of substituted energy sources and their associated emissions, the increase in Norwegian domestic emissions from production, and how these variables evolve over time. Rystad and Vista have, for example, used different estimates for demand elasticity. Based on a review of 11 studies, Rystad estimates the demand elasticity to be -0.11. Vista uses the result from a meta-study of 75 research papers that estimate a demand elasticity of -0.26. Different choices of parameters like this will give different conclusions. The ongoing debate underscores the importance of considering the climate impact of the Norwegian oil industry, providing additional motivation to pursue our research question.

Ahlvik et al. (2022) also examine supply-side policies in the oil market. They discover that companies respond to taxes by reducing oil exploration, specifically that a one percentage point higher royalty rate decreases oil exploration by 3%. These taxes increase the oil price, which has long-term distributional impacts on the oil market. The authors observe that the surplus is moved from consumers to producers and governments, compared to a case where all production-based taxes are removed (Ahlvik et al., 2022). They also analyse a hypothetical supply-side policy regarding a climate royalty on new oil discoveries. Their results indicate that if a single price-taking country were to apply it, such a policy would only have an effect of 9–20% of the intended emission reduction. This is because of oil market leakage due to increased supply as a result of higher oil prices. Further, they argue that this leakage could be avoided if the royalty is adopted globally. The authors discovered that emission was reduced almost linearly when a global climate royalty surcharge is set beyond today's tax level. They estimated a reduction of around 0.16 GtCO2 per percentage-point increase in the rate (Ahlvik et al., 2022). This essentially means that if one country reduces its oil supply, the emission reduction is very small, and we would need a global agreement on supply to get the desired effect.

The paper by Ahlvik et al. (2022) also comments on short and long-run supply elasticities. The authors find that the oil supply is very inelastic in the short run, which is consistent with previous literature, i.e., Baumeister and Hamilton (2019) who found the short-run oil supply elasticity to be 0.15. However, given the nature of climate change, the long-run elasticity is more important. The paper looks at oil exploration and estimates a long-term supply elasticity of 1.96, which is higher than previous literature typically suggests (Ahlvik et al., 2022). The paper only takes into account the direct leakage through the oil market. A likely reason for the higher estimate is that they look at oil exploration and not production. A second one is that markets may respond more to taxes than price changes if taxes is seen as more permanent in the long run.

2.3 Supply vs. Demand

Before going any further, it will be useful to look at what the literature on oil markets views as the main drivers of oil price fluctuations. The work by Hamilton (1983; 1985) speaks in favor of supply shocks being the biggest driver of oil prices. Kilian (2009) later challenged these findings and used a structural VAR method to argue in favor of the demand side. The debate is still ongoing and highly relevant.

In his influential work in 2009, Lutz Kilian made significant contributions to understanding oil prices and the role of demand and supply shocks. Kilian's research emphasized the importance of distinguishing between different types of shocks, particularly demand and supply shocks, in explaining oil price movements. He argued that demand shocks, stemming from shifts in global economic activity and consumer behavior, played a substantial role in driving fluctuations in oil prices. Kilian's research on oil price shocks between 1973 and 2007 reveals that supply disturbances were generally short-lived. However, the impact of demand shocks, including changes in global activity and expectations of future oil supply, was significant. Kilian argues that even during physical supply disruptions, it is the precautionary demand component that primarily drives price increases, rather than the supply shock itself. Importantly, Kilian's results contradict Hamilton's perspective, suggesting that the supply side is less important.

Hamilton's view on oil supply shocks centers around the idea that exogenous factors, such as wars and conflicts in oil-producing regions, play a significant role in causing fluctuations in oil prices. In his research papers from 1983 and 1985, he sought to explain why oil prices consistently spiked just before U.S. recessions based on post-war data. According to Hamilton, these price surges were not primarily driven by domestic economic conditions or the U.S. business cycle. Instead, he argued that external disruptions to oil supply, particularly in regions with significant oil production, were the primary drivers of these price fluctuations. In a more recent paper, Baumeister and Hamilton (2019) revisit the role of oil supply and demand shocks in the context of VAR models. They recognize that traditional identification strategies often rely on strong assumptions, which may lead to biased estimates and potentially misleading interpretations. To address this concern, they use a Bayesian approach that incorporates insights from economic theory and allows for a more robust identification of shocks. With this framework, they revisit earlier studies on oil supply and demand and discover supporting evidence for the role of both factors in influencing oil prices, however they conclude that supply factors seem to be more important. They also find that supply shocks in oil markets lead to a subsequent reduction in economic activity, while price increases due to increased oil consumption demand have minimal impact on economic activity.

The literature on oil markets offers divergent perspectives on the main drivers of oil price fluctuations. Ultimately, a comprehensive understanding of oil price dynamics requires considering both supply and demand factors and the dynamic interplay between the macroeconomy and the oil market. The ongoing debate highlights the complexities of the oil market and the need for nuanced analyses.

2.4 Temporary vs. permanent supply shocks

This paper examines the effects of both a permanent and a transitory supply shock on Norwegian oil production. It was recently discussed by Rebei and Sbia (2021) that these shocks have different effects on global oil market variables. Knowledge about this is important to interpret our findings.

Permanent oil supply shocks are sustained shifts in oil supply that persists over a long period and is hard to reverse. This could be due to a variety of factors, such as fundamental changes in production, technological advancements, geopolitical events, and so on. An example of a historical event resulting in a persistent supply shock was the Iranian Revolution of 1979. The revolution caused a significant reduction in oil production, accounting for 7% of world production at that time (Federal Reserve History, 2013). Consequently, oil prices started to rise in mid-1979, and by February 1981, they reached a local all-time high of $39 per barrel, up from around $16 per barrel before the revolution. It is worth noting that the supply disruption resulting from the Iranian Revolution may not have been the only cause of the rising oil prices. Some of the rise can also be attributed to the fear of further disruptions and speculative hoarding (Federal Reserve History, 2013). Nevertheless, the event did stimulate non-OPEC oil production, with increased investments in exploration and production, which helped bring oil prices back down. However, global production did not fully recover until 1990, a decade after the Iranian Revolution.

Transitory oil supply shocks refer to temporary oil supply shifts expected to be resolved relatively quickly. These shocks are usually caused by temporary disruptions such as geopolitical tensions, weather events, or refinery outages. An instance from history that illustrates a transitory supply shock is the reduction in oil production following Hurricane Katrina in 2005. The severe damage caused by the hurricane to U.S. refineries led to an increase in oil prices by approximately 7%, according to our data on oil prices. The shock classifies as transitory because it does not change the fundamental situation in the oil market over a longer period. Then-sitting President George W. Bush intervened by releasing 30 million gallons of oil from the Strategic Petroleum Reserve (SPR), which helped stabilize the prices (CFR, 2005). Another example of a temporary event is the 1986 shock to Norwegian production, which we will delve into in Section 3.

To investigate how the oil price and oil availability react to a permanent vs. transitory shock, Rebei and Sbia (2021) applied a vector autoregressive model with unobserved components. Their findings indicate that both transitory and permanent oil supply shocks have similar qualitative effects on the variables, but the magnitude of the responses and the timing of the highest impact can differ between the two types of shocks. By studying a one percent increase in the structural shock to oil production, they found that a transitory shock tends to generate more significant contemporaneous effects on oil availability than a permanent shock. As oil availability increases more on impact, we also see larger contemporaneous effects on the oil price in the temporary case. Further, oil availability response fades within eight quarters for a transitory shock compared to a permanent shock where the response is persistent. The persistence of oil price effects is comparable for both types of shocks, but transitory shocks seem to generate longer-lasting effects. Lastly, their findings suggest that permanent shocks exert a more prolonged impact on global output, resulting in a larger cumulative response (Rebei & Sbia, 2021).

We find Rebei and Sbia's results valuable but want to question their results regarding persistence in oil prices. We believe that permanent oil supply shocks will have a longer-lasting impact on oil prices than temporary shocks. This is because when there is a permanent supply shock, oil agents know that the increase in oil availability will be sustained, so they adjust their behavior more gradually, and the impact on prices is less immediate. In contrast, when there is a temporary supply shock, agents know that the increased availability will be short-lived, so they adjust their behavior quickly, driving prices back up faster. It is also worth mentioning that Rebei and Sbia's results show that oil prices recovered with the same amount (7%) for both shocks, implying that the difference in persistence originates from the contemporaneous effects on oil availability.

Another plausible reason for the differences in oil price persistence is effects through changed output. Gross domestic product (GDP) affects oil prices through its impact on demand for oil. Increased GDP stimulates economic activity, leading to increased demand for energy, including oil. As demand for oil rises, prices tend to increase, resulting in a faster recovery for oil prices.

3 1986 Oil Strike

We can analyse historical exogenous shocks to Norwegian oil production to investigate the effect of a shutdown in the Norwegian sector. Such an event study will enhance the interpretability of our model findings and make it easier to judge their validity.

On April 6th, 1986, a strike by kitchen workers forced a temporary shutdown of several Norwegian oil rigs (Meland, 2020). The strike resulted from failed wage negotiations between the catering workers' union and the employer's federation of the Norwegian oil industry. Around 3500 workers were evacuated from the oil rigs, resulting in a shutdown of 62% of Norwegian oil production (Caldara et al., 2019). At that time, Norway produced 900,000 barrels daily, making up 1.5% of the world's oil output (Lyngve, 1986). The cut in Norwegian production resulted in a 0.97% decrease in global oil supply (Caldara et al., 2019). On April 25th, the Brundtland III Cabinet decided the case should go to the compulsory wage board[3], and oil production was immediately restored (Meland, 2020).

The strike, an unexpected event from outside the oil industry, caused an exogenous shock to Norwegian oil production. Caldara et al.'s paper from 2019 also classifies the shock as exogenous, making it a useful event study for investigating the impact of a Norwegian oil supply shock. To assess the impact of the shock on the oil market, we first had to determine the timing at which the market became aware of the kitchen workers' strike. A comprehensive search in newspapers between March 20th and April of 1986 was conducted to determine the timing of the news coverage. We used Atekst (2023) and Nasjonalbiblioteket (2023)

[3]The compulsory wage board is a government-established body that intervenes in negotiations between the parties in a labor dispute if a strike is considered to be a danger to "vital societal interests".

to find newspapers published in Norway, and Newspaper.com (2023) for newspapers published in the U.S. The search focused on papers containing the keywords "oil," "strike," and "Norway." Relevant newspapers were manually checked, counted, and allocated to their respective dates. Doing this gave us an image of when the national and international markets became aware that Norwegian oil production would be halted. The results, along with the oil price between March 14th and April 15th, are plotted in Figure 1.

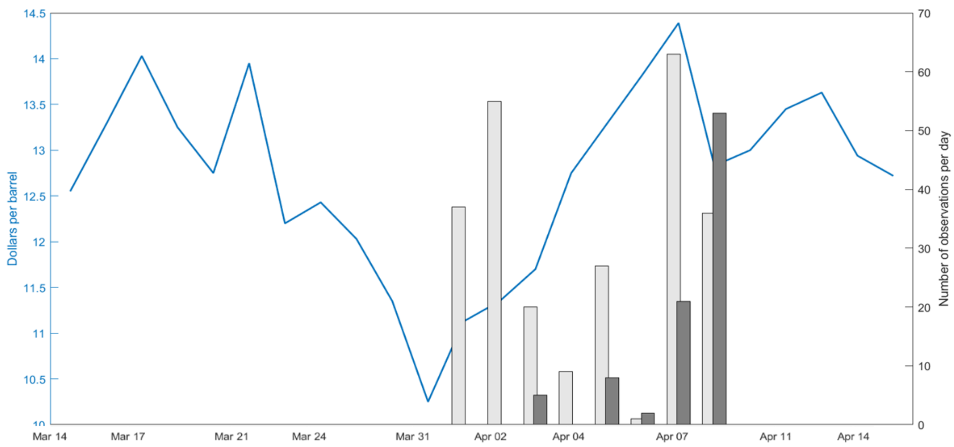

Figure 1: The oil price is displayed as a line with corresponding values on the left Y-axis. The light gray bars indicate relevant Norwegian newspapers, while the dark gray ones indicate American newspapers. The right Y-axis shows the corresponding numbers. We stopped counting newspapers after April 8th because, at this point, the market was well aware of the strike. 14.03.86–15.04.86, daily frequency.

News about the oil strike first surfaced in Norway on April 1st and subsequently reached the international market with a slight time lag. A New York Times article from April 7th, 1986, writes, "When the possibility of the strike became apparent last Friday, the price of Brent crude oil, a benchmark grade, advanced to $12.20 a barrel, up $1.20 from the previous day's price" (Lohr, 1986), which indicates that the market was not aware of the strike before early April. The graph also illustrates a significant 27% increase in oil prices between April 1st and 7th, going from \$11.30 to \$14.35. As we can see, the increase in oil prices is consistent with the period when news about the strike came out. However, we cannot say with certainty that the entire 27% increase is due to news about the Norwegian production cut. The oil price was already at a local minimum on April 1st, largely due to a failed OPEC meeting on March 24th (Woldsdal, 1986). OPEC ministers met to agree on production quotas to push the price up, but when the meeting collapsed, the price declined instead. Recovery from this, and natural fluctuations, may also have contributed to the price rise. However, to our knowledge, there were no other significant happenings between April 1st and 7th that could influence the oil price, making it reasonable to believe that news about the Norwegian strike played a significant role. A United Press International news article from April 25th, 1986, tells us that the oil price also responded to the strike ending: "North Sea oil prices nosedived Friday as Norwegian oil workers ended a 21-day strike that has removed 900,000 barrels a day of oil from the glutted world market" (Liston, 1986), which is further evidence that the oil price responds to Norwegian production.

There were no significant changes caused by the strike regarding global oil production. According to data obtained from the Energy Information Administration (EIA), there was a minimal increase in global production from 55,162 TBPD to 55,266 TBPD between April and May 1986. Conversely, Norwegian production declined from 861 TBPD to 324 TBPD during the same period. Given that Norwegian production is a component of global production, it can be argued that since global production did not decrease, production must have been augmented elsewhere. If this supposition holds, it becomes evident that other suppliers intensified their production to compensate for the demise of Norwegian production. Another explanation for the absence of changes in global production could be attributed to natural fluctuations in production.

Despite the fact that the strike resulted in a 62% reduction in national production, its impact on global production amounted to only 0.97%, as stated previously. Monthly variations in global production exceeding 1% are not uncommon and can be responsible for the rebound in global production. This suggests that other suppliers did not respond to the demise of Norwegian oil production. It is hard to determine whether or not the strike had a beneficial effect on net emissions without knowing the reason for unchanged global production. If the first stated reason is true, net emissions will remain unchanged as other suppliers would fully compensate for the reduction. On the other hand, if the latter reason holds true, the strike had a beneficial effect on net emissions, but this effect is hard to observe in the raw data due to fluctuations. To answer our research question, it is important to identify the true reason for the rebounded production.

According to The New York Times (Lohr, 1986), the oil market perceived the strike as temporary, without affecting the fundamental situation of the oil market. To further ascertain the veracity of this claim, we have analysed certain key indicators in the West Texas Intermediate (WTI) futures market. If the market perceives the production stop as temporary and expects the supply to return to normal shortly, we might observe a backwardation[4] in the futures market. The behavior of short-term[5] futures contracts can also reveal insight into the market's belief regarding the strike's duration. If the prices of near-term futures contracts are relatively stable or only experience a temporary increase, it may suggest that the market believes the production strike to be a short-term event.

[4]Backwardation: When the spot price of a commodity is higher than the futures price.

[5]Futures contracts that have a duration of up to the 12th position are classified as short-term or near-term.

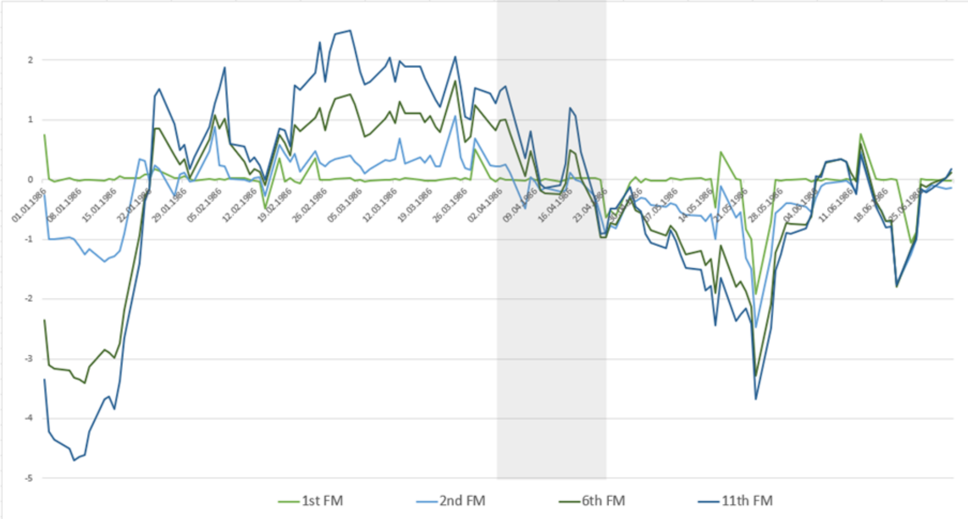

Figure 2: Plot of the 1st, 2nd, 6th, and 11th position of the NYMEX WTI futures prices in deviations from the WTI spot price for Crude oil. 01.01.86–25.06.86, daily frequency. The shaded area marks the strike. Y-axis: Deviations from spot in dollars.

Through the computation of the difference between the WTI spot price and futures prices across different positions, an analysis of their relationship surrounding the strike was conducted, as illustrated in Figure 2. The figure shows that during the period between the 22nd of January and the 20th of April, the market was predominantly in a state of contango[6], and in backwardation for the rest of the sample. Notably, the data reveals an intriguing trend of the market converging from contango to backwardation during the strike, particularly evident in the 6th and 11th positions.[7] This phenomenon can be attributed to the realization in the market that the strike would be of a temporary nature. The precipitous increase on April 16th was a consequence of a swift decrease in the spot price where the 6th and 11th position futures contracts demonstrated a comparatively tardy reaction as opposed to the 1st and 2nd position. The near-time futures contracts also look

[6]Contango: When the futures price of a commodity is higher than the spot price.
[7]The 11th position was chosen due to lack of data on the 12th position future contracts from the date 01.07.1986–31.07.1986.

relatively stable without any notable deviations during the strike (see Appendix A.1), strengthening the argument that the market viewed the strike as a temporary event.

Previous literature (see Baumeister and Hamilton 2019, Caldara et al. 2019) argues that the short-run oil supply is very inelastic. This is also the case in April 1986 where the actual price elasticity of supply is calculated to be 0.087.[8] Therefore, given that the market believed the strike to be temporary, it is assumable that oil suppliers did not react to the shock and that the lack of change in global production is due to natural fluctuations.

In conclusion, our inquiry into the oil market pertaining to the Norwegian strike has yielded valuable insights. Our findings suggest that the shock had a significant impact on the price of oil. The oil price is known to be highly volatile and responsive to market changes and news. Over the 7-day period when news about the Norwegian strike surfaced, the price rose 27%. Total global production did likely fall as Norwegian production fell. This is however hard to conclude with certainty, as we cannot observe what global production would have been in April 1986 if the strike had not happened. By reading newspapers from 1986 and investigating the WTI futures market, we found that the market believed the strike to be temporary. It is, therefore important to state that a more persistent shock of greater magnitude may induce a reaction from other oil-producing countries, dampening the effect on global production and hence net emission reduction. Lastly, it's also important to acknowledge that today's oil market may differ structurally from that of 1986, which underscores the need to assess the external validity of

[8] $e_p = \frac{\%-change\ in\ production}{\%-change\ in\ price} = \frac{0.0214}{0.247} = 0.087$. % - change in production is the change in global oil production from 1st of April 1986 to 1st of May 1986. % - change in price is the change in the real price of oil from 1st of April 1986 to 1st of May 1986.

our findings carefully. Any attempt to generalize these findings beyond the scope of our analysis must be approached with caution to avoid misleading conclusions.

4 Methodology

Through an examination of the impact of the Norwegian strike on global production and prices, we have gained insight into the characteristics of a temporary shock to Norwegian production. However, these findings were based on the global context of 1986, necessitating further investigation to determine if the same results persist into the present day. To achieve this objective, we will employ a structural vector autoregression (SVAR) model that reproduces the size and magnitude of the 1986 strike. If the findings are consistent under the present-day conditions, they can be used to interpret our final results.

To comprehensively address our research question, we must generate a permanent shock to Norwegian production. To accomplish this, we will apply the conditional forecasting approach described in Section 4.2.

4.1 Structural Vector Autoregression (SVAR) approach

There is no explicit agreement among researchers on how to model global oil markets. However, using Vector Autoregression (VAR) models combined with identification strategies and explanatory variables has become increasingly popular in the field since Kilian's (2009) seminal paper. This approach builds on Sims' (1980) work and is based on a time when the validity of traditional large-scale dynamic simultaneous equation models and exogeneity assumptions were questioned. Structural VAR models have appealing properties, such as the ability to generate impulse response functions and variance decompositions, which allow for causal inference and assessment of the relative impact of exogenous shocks on variables. However, whether these models provide true causal inference relies heavily on the credibility of the identifying assumptions being imposed.

4.1.1 The general SVAR setup and identification strategy

A general SVAR(p) model with n variables can be written as Equation 1. B_0 is the impact matrix containing $n \times n$ elements and shows how the variables of interest respond to shocks ε_t at the current date t while Y_t is a vector of endogenous variables, and b is a vector of constants. B_j captures all autoregressive coefficients and has the same dimensional form as B_0. ε_t is a vector of uncorrelated structural shocks, each with unit variance. All vectors contain $n \times 1$ elements.

$$B_0 Y_t = b + \sum_{j=1}^{p} B_j Y_{t-j} + \varepsilon_t, \quad \text{where } \varepsilon_t \overset{iid}{\sim} N(0, I) \tag{1}$$

The issue is that this model cannot be estimated as the elements of B_0 cannot be identified. But if B_0 is invertible, we can pre-multiply on both sides to obtain the reduced form VAR shown in Equation 2.

$$Y_t = a + \sum_{j=1}^{p} A_j Y_{t-j} + e_t, \quad \text{where } e_t \overset{iid}{\sim} N(0, \Sigma_e) \tag{2}$$

Where $a = B_0^{-1} b$ is a vector of constants, $A_j = B_0^{-1} B_j$ is a $n \times n$ matrix, and $e_t = B_0^{-1} \varepsilon_t$ is a vector of error terms. The covariance matrix Σ_e is positive semi-definite, and symmetric, often expressed as Equation 3.

$$\Sigma_e = \mathbb{E}[e_t e_t^{'}] = B_0^{-1} (B_0^{-1})^{'} \tag{3}$$

By making use of a lag operator, we can rewrite the reduced VAR as:

$$A(L) Y_t = a + e_t, \quad \text{where } A(L) = (I - \textstyle\sum_{p=1}^{p} A_p L^p)$$

Pre-multiplying this with $A(L)^{-1}$ gives us the moving average (MA) representation of the reduced form VAR:

$$Y_t = v + B(L)e_t, \quad \text{where } \ B(L) = A(L)^{-1} \ \text{ and } \ A(L)^{-1}a = v$$

The inverse of $A(L)$ exists if all its eigenvalues are less than unity in absolute value, and the model is then stable. Given this, this reduced form VAR can be estimated using standard methods. In this system, a single shock is unlikely to occur in isolation, as the covariance matrix of e_t is not diagonal and the errors can be correlated. This makes it challenging to infer causality. We can express the reduced form errors using a linear combination of a matrix that outlines the structural connections between the uncorrelated (structural) shocks. We denote $e_t = B_0^{-1}\varepsilon_t$. However, we still have the same problem as before, as we have an underidentified system of equations. The SVAR has n^2 contemporaneous parameters, while the reduced form VAR has $n(n+1)/2 < n^2$ parameters in the covariance matrix. Since Σ_e is a positive definite symmetric matrix, we can achieve exact identification by restricting $B_0^{-1} = S$, where S is the lower Cholesky factor of Σ_e. I.e., we impose exclusion restrictions on B_0^{-1}. By assuming that $\mathbb{E}[e_t e_t^{'}] = I$.[9] We then get:

$$Y_t = v + \Theta(L)\varepsilon_t, \qquad \text{where} \quad \Theta(L) = B(L)S$$

By identifying S, we can compute $\Theta(L)$ through the reduced form $B(L)$. For the model to give wanted results, we must be able to identify the structural parameters in Θ_j.

The original solution to the identification problem suggested by Sims (1980) was to take in use a recursive system, i.e., restrict all elements of B_0^{-1} above the main diagonal to be zero. This is possible for two reasons. Firstly, there is exactly $n(n-1)/2$ elements above the diagonal. Secondly, we can make use of Cholesky Decomposition. However, by

[9]If S is lower triangular, we can write: $\mathbb{E}[e_t e_t^{'}] = S^{-1}\,\mathbb{E}[e_t e_t^{'}](S^{-1})^{'} = S^{-1}(SS^{'})(S^{-1})^{'} = I$.

making use of this identification strategy we are imposing assumptions on the contemporaneous (structural) relationships between the shocks and the variables in Y_t. As Θ_0 (The structural MA coefficients for the companion form at time $t = 0$) only contains zeros in the upper triangular, variables in Y_t ordered below do not affect variables ordered above contemporaneously, making the ordering of variables essential. This assumption does not need to be fulfilled in the following periods ($j > 0$) as $\Theta_{j(j \neq 0)}$ are full matrices not subject to any restrictions. Lastly, to derive the long-run impacts of a shock, we compute impulse response functions (IRFs) expressed in Equation 4.

$$\Theta(1) = \sum_{h=0}^{\infty} \frac{\partial Y_{l,t+h}}{\partial \varepsilon_{i,t}} = \theta_{li,h}, \quad \text{where } l = 1, \ldots, n \text{ and } i = 1, \ldots, n \tag{4}$$

4.1.2 SVAR model including Norwegian production

Our starting point builds upon Kilian's (2009) previous work. Kilian utilized three variables at monthly frequencies: global crude oil production, a self-constructed index of real activity, and the real price of oil. Our specification aligns with Kilian's approach, with two key distinctions. Firstly, we expanded the model by including Norwegian production to evaluate the significance of Norwegian crude oil production in the oil market. Secondly, we have modified the measure for economic activity by employing the monthly world industrial production (WIP) index developed and maintained by Baumeister and Hamilton (2019). The reason for this is that Kilian's index has already been transformed using the entire sample period, implying that forecasting results might be subject to a look-ahead bias[10] (Baumeister & Guérin, 2021, p. 7). As we want to be consistent in the choice of index, and we want the index

[10]Look-ahead bias can arise when a study or simulation incorporates information or data that was not accessible or known during the time frame under analysis.

to reflect the current economic conditions, the WIP index is chosen for both models. For more information about our data, see Appendix A.2.

$$\begin{bmatrix} \Delta nor.prod \\ \Delta g.prod \\ rea \\ lrpo \end{bmatrix}_t = \begin{bmatrix} \theta_{11} & 0 & 0 & 0 \\ \theta_{21} & \theta_{22} & 0 & 0 \\ \theta_{31} & \theta_{32} & \theta_{33} & 0 \\ \theta_{41} & \theta_{42} & \theta_{43} & \theta_{44} \end{bmatrix} \begin{bmatrix} \varepsilon^{\Delta nor.prod} \\ \varepsilon^{\Delta g.prod} \\ \varepsilon^{rea} \\ \varepsilon^{lrpo} \end{bmatrix}_t + \Theta_1 \varepsilon_{t-1} + \cdots + \Theta_p \varepsilon_{t-p}$$

Where $\Theta_i, (i = 1, 2, \ldots, p)$ are full matrices, and p is the number of lags.

The model presented above follows the ordering of variables as outlined by Kilian (2009). Supply variables are at the top, followed by global demand, and finally the oil price. The chosen ordering in the SVAR model reflects the causal relationships and economic logic associated with the variables and the oil market. Norwegian oil supply is placed first, followed by global oil supply. Supply variables preceding demand, implies that oil producers do not adjust their production levels in response to shocks in aggregate demand within a month. This exclusion restriction is plausible because oil producers are typically slow to respond to demand shocks due to high adjustment costs. Supply-side factors are typically considered to have a more immediate impact on the market compared to demand-side factors. Supply shocks can quickly affect the availability of oil in the market, leading to changes in prices and subsequent adjustments in demand. Finally, the oil price variable is placed last as it is influenced by both supply and demand factors. Although oil prices can be observed daily, economic agents take time to adjust their behavior, resulting in a delayed impact on the level of real economic activity.

To ensure the proper transmission of oil price shocks, considering multiple time lags is crucial. We use monthly data from 1974:M01 to 2019:M12, making it unproblematic to add many lags. Following the advice of Hamilton and Herrera (2004), we specify a high lag order and set it to 24 to mirror Kilian (2009). This implies two years of lagged endogenous variables.

4.1.3 Results from the SVAR model

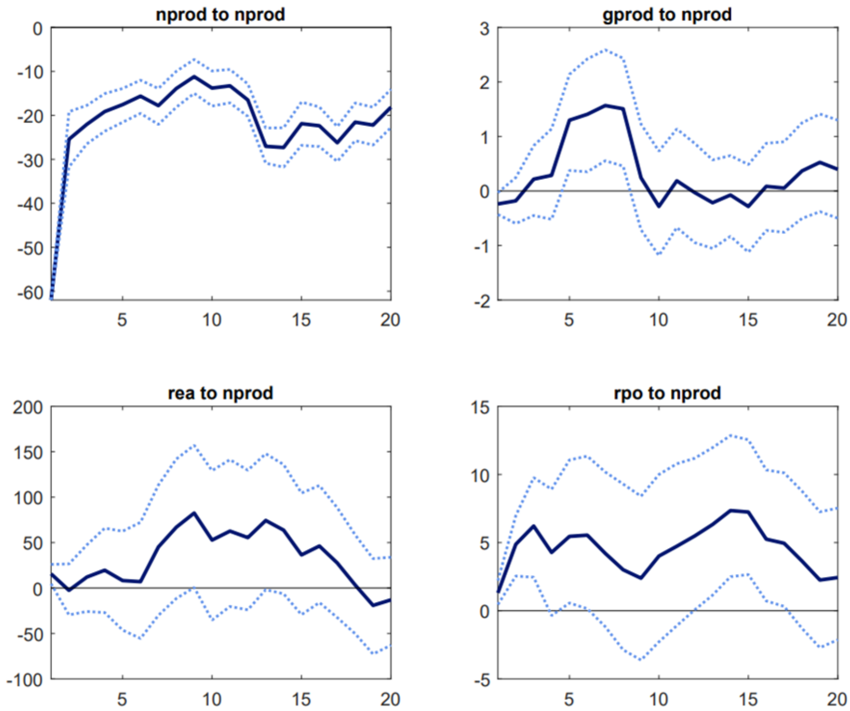

Figure 3: Impulse responses of a 62% temporary reduction in Norwegian oil production. Nprod and gprod are in accumulated responses. The oil price is deflated and logged. Confidence Level: 68%.

The impulse response functions show how global oil production, oil price, and real economic activity respond to a 62% negative temporary shock to Norwegian oil production. We have shocked Norwegian production by 62% to mimic the shock in conjunction with the 1986 strike to see if our model produces the same results as we saw in 1986. Impulse responses that are not accumulated are included in Appendix A.3.

The biggest effects are seen in the price. The sudden decrease in oil production makes the price increase. The price peaks after 3 months at a 6.5% increase, and the effects are significant with 68% confidence bands for about six periods. In contrast, during the 1986 strike, which also involved a temporary exogenous shock, the price increase was 27%. However, as explained in Section 3, there might have been other factors contributing to this rise. While the direction and magnitude of the shock are not far away from the findings of Rebei and Sbia (2021), the persistence differs. Their VAR model estimates that a one percent[11] temporary positive oil supply shock results in an immediate price decline of -9.5% and that the price remains significantly below its long-term level for more than five years. Our model only shows significant price change for around 6 months. As mentioned in Section 2.3, the literature extensively debates the significance of exogenous oil supply shocks in explaining the fluctuations observed in oil prices. Our model reports that the oil price reacts to supply shocks, which aligns with Hamilton (2019) and the supply-side view.

Our model reports a short-term lagged increase in accumulated global production in response to the shock. The increase is significant from period five to eight and lies around 1.5%. The delayed response may be attributed to the short-term inelasticity of the oil market. It is interesting to see that other producers change their behavior within such a relatively short period. One plausible explanation for the increased production is the rise in oil prices accompanying the shock. The impulse response function shows that other oil-producing countries produce more in a short period after the shock, before then producing less, resulting in no significant accumulated change in the long-term. Considering that global accumulated production remains unchanged in the long run and Norwegian accumulated production stabilizes at a lower level, this shock

[11] 62% of Norwegian oil supply equals 1.32% of global oil supply in our data, so the shocks in Rebei and Sbia (2021) are a bit smaller on aggregate (1% vs. 1.32%).

would reduce total oil production, which means that the shock would lead to lower CO2 emissions from oil. The same applies for the 1986 strike where Norwegian production returned to its original level after around one month, and accumulated production stabilizes at a lower level due to the disruption of the strike.

One thing to note is that the negative shock of 62% in our model is just an inverse linear scaling of a positive 1% shock. In other words, we have assumed linearity. It is important to be aware of the limitations of our methodological approach. In econometric modeling, shocks with very large magnitudes can have nonlinear effects on the variables being analysed. It is possible that the linear scaling of the shock might not accurately capture the full dynamics and complexities of the system. As a result, the effects observed in our SVAR model may not fully represent the real-world response. It is also worth noting that the effect on price is not significant at a 95% confidence level.

4.2 Conditional Forecast approach

To create a permanent shock to Norwegian oil production, we apply a conditional forecast. Waggoner and Zha (1999) provide a conditional forecast approach which is described in "Applied Bayesian econometrics for central bankers" written by Andrew Blake and Haroon Mumtaz (Blake, & Mumtaz, 2017). This is a handbook for key topics in Bayesian econometrics from an applied perspective. The approach described is convenient for two reasons. First, it allows us to choose the duration and size of the shock to Norwegian production. Second, it allows us to analyse the response from global oil production and the real price of oil.

4.2.1 The conditional forecast approach in our context

First, assume a standard VAR(4) model iterated K times forward, as shown in Equation 5.

$$Y_{t+K} = c\sum_{j=0}^{K} B^j + \sum_{j=1}^{4} B^j Y_{t-1} + A_0 \sum_{j=0}^{K} B^j \varepsilon_{t+K-j} \tag{5}$$

Y_{t+K} is the K period ahead forecast and has been decomposed into three components. When we constrain the J^{th} variable in Y_{t+K} to follow a fixed path, we do also restrict the future structural shocks (ε_{t+K-j}). Waggoner and Zha (1999) express these constraints on future innovations as

$$R\varepsilon = r \tag{6}$$

Where r is a $(M \times k) \times 1$ vector. M represents the number of constrained variables and k denotes the number of periods. In our model, only Norwegian oil production is constrained ($M = 1$), and the forecast horizon is set to five years ($k = 60$).[12] R is a matrix with dimensions $(M \times k) \times (N \times k)$ where N is the number of variables in Y_{t+K}. The $(N \times k) \times 1$ vector ε contains the constrained future shocks and has a least square solution that can be represented as $\hat{\varepsilon} = R'(R'R)^{-1}r$ (see, Doan et al. 1983).

Further, r and R can be written as shown under. $\hat{\varepsilon}$ is stacked as in Equation 7.

$$r = \begin{pmatrix} P - \tilde{X}_{t+1} \\ P - \tilde{X}_{t+2} \\ \vdots \\ P - \tilde{X}_{t+k} \end{pmatrix} \quad , \quad R = \begin{pmatrix} z^1_{1,1} \cdots z^N_{1,1} & 0 & \cdots & 0 \\ z^1_{2,1} \cdots z^N_{2,1} & z^1_{1,1} \cdots z^N_{1,1} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ z^1_{k,1} \cdots z^N_{k,1} & z^1_{k-1,1} \cdots z^N_{k-1,1} & \cdots & z^1_{1,1} \cdots z^N_{1,1} \end{pmatrix}$$

Note that P represents the values of the constrained path chosen, and $\tilde{X}_{t+1}, \tilde{X}_{t+2}, \ldots, \tilde{X}_{t+k}$ represents the unconditional forecasts for the constrained variable at all k periods (See Appendix, A.4). P is set to zero for all 60 periods to shut down Norwegian production. The R matrix shows all elements in the impulse responses made by the constrained variables to the structural shocks ε at horizon $1, 2 \ldots k$. In our model, $\hat{\varepsilon}$ has 240 rows. This implies that R has to be structured such that it contains (60×240) elements for Equation 6 to hold.

$$\hat{\varepsilon} = \begin{pmatrix} \hat{\varepsilon}_{1,t+1} \\ \vdots \\ \hat{\varepsilon}_{N,t+1} \\ \vdots \\ \hat{\varepsilon}_{1,t+h} \\ \vdots \\ \hat{\varepsilon}_{N,t+h} \end{pmatrix} \tag{7}$$

[12] Horizon is 60 periods as the data is at monthly frequency.

After $\hat{\varepsilon}$ is calculated using the least square formula, we reshaped it to a (60×4) matrix separating the shocks in different columns. We calculated the conditional forecast in the same way as the unconditional forecast. The only difference is that the conditional approach accounts for the restricted shocks.

Waggoner and Zha (1999) introduced a Gibbs sampling algorithm[13] that we employed to determine the distribution of the conditional forecast. They demonstrated that it is possible to compute both the mean ($\bar{M}$) and variance ($\bar{V}$) of the distribution of the restricted future shocks ε by applying the formulas shown in Equation 8. A detailed explanation of the sampling algorithm is included in Appendix A.5.

$$\bar{M} = R'(R'R)^{-1}r \ , \ \ \bar{V} = I - R'(RR')^{-1}R \tag{8}$$

We iterated the algorithm for 10,000[14] $REPS$ and set a burn-in period of 2,000 iterations. This implies that the first 2,000 iterations are not accounted for in the final forecast to try to make sure that the sample has converged. When the algorithm is done, the conditional forecasts are stored in four different outcome matrices, one for each variable. The outcome matrices in our model include $8,000 \times 587$ elements, where the number of rows reflects the number of iterations ($REPS - BURN$), and the columns represent both the past and forecasted data. The first 527 columns represent the past data, where the first 24 periods are excluded because of the lags. The last 60 columns are the posterior distribution at each forecast horizon, where the median of each column

[13]Gibbs Sampling is an iterative Monte Carlo Markov Chain technique employed to estimate intricate joint distributions. It operates by drawing instances from the distribution of each variable, conditioned on the current values of the remaining variables (Makin, 2020).

[14]Increasing the number of draws above 10,000 comes at a great computational cost and earlier attempts suggest that the sampler converges very quickly making more draws unnecessary to perform valid inference.

creates the median forecast shown in Figure 4 (Blue line). Examples of these distributions are plotted and shown in Appendix A.6.

4.2.2 Results from the conditional forecast approach

When plotting the forecast conditional on a shutdown of Norwegian production, we obtain the results represented in Figure 4 (See, Appendix A.7, Figure 16 for zoomed-out plot). These results argue that global production will increase for the first eight months when looking at the median forecast (50th percentile). Thereafter, we see production decreasing and stabilizing at a lower level. The price of oil increases for over a year before it flattens out towards the end of the horizon. Real economic activity shows a large increase for the first part of the forecast before decreasing back to its initial value.

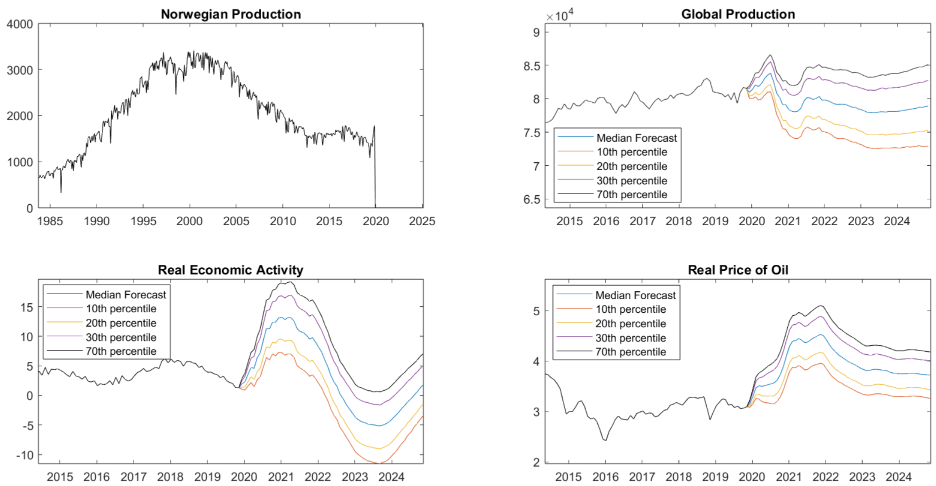

Figure 4: Plot of CF results in the zero-restricted scenario. Dataset: 1974:M2 - 2019:M12. Forecast horizon: 60 periods (five years). Y-axis: TBPD.

The forecast is done on raw data, and we see an upward trend in global production. To account for how this affects our results, we made a forecast where Norwegian production is restricted to 1,476[15] TBPD, which

[15]We change the path by restricting P in the r vector to 1,476.

is the mean production of the last 24 observations.[16] The responses from the other variables will work as a picture of what would happen if Norwegian production was unchanged; see Appendix A.7, Figure 17. At first glance, the responses from all variables look flatter. However, to compare the two states more easily, we have plotted the median forecast for both the zero-restricted and the mean-restricted scenarios in Figure 5.

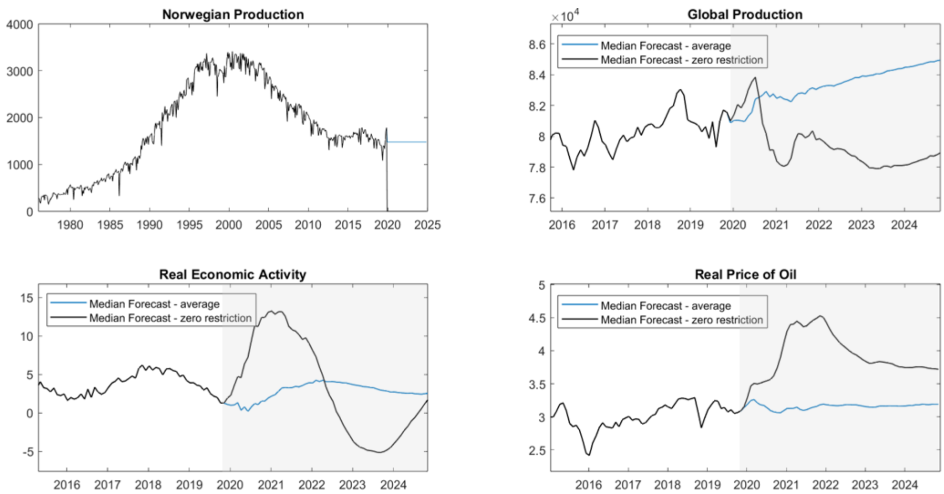

Figure 5: Comparison of the median forecast for both the mean-restricted scenario (Blue line) and zero-restricted scenario (Black line). Y-axis: TBPD.

Global production in the zero-restricted forecast lies above the mean-restricted forecast for the first ten months, indicating that other producers adjust their production in action to the demise of Norwegian production. These results align with the short-term results obtained in the structural VAR model. In both the 1986 strike and the structural VAR representation, we concluded that global production was unchanged in the long-term and hence, had some beneficial effects on the aggregated CO2 level. To make arguments regarding the impact on the climate

[16]We evaluate the mean of the last two years as a good estimate of future unrestricted oil production.

from this permanent shock, we must consider the total amount of barrels produced. Since our global production variable excludes Norwegian production, the blue line in Figure 6 shows the total amount of barrels produced each day in the mean-restricted scenario.

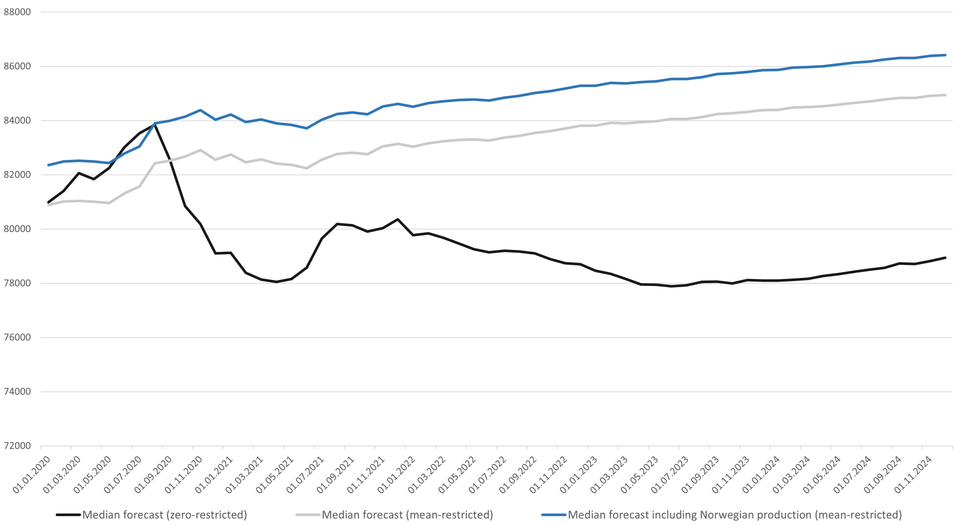

Figure 6: Plot of CF results. Black line: Total oil production in the zero-restricted scenario. Blue line: Total oil production in the mean-restricted scenario. Gray line: Global oil production in the mean-restricted scenario (excluding Norwegian production). Y-axis: TBPD.

By comparing the mean-restricted scenario (blue line, Figure 6) with the zero-restricted scenario (black line, Figure 6), we see that shutting down Norwegian oil production has no effect on CO2 emissions for the first nine months as other producers will increase their production in response. However, from a longer perspective, shutting down Norwegian production results in fewer barrels produced each day from the ninth month until the end of our horizon. It is evident that global production has consistently decreased, leading to a smaller overall cumulative production in the zero-restricted scenario. The new level is approximately 7.2% lower than production would have been, which we consider to be a unrealistic decline, especially considering that the Norwegian production shutdown only accounts for a 2% reduction in global production.

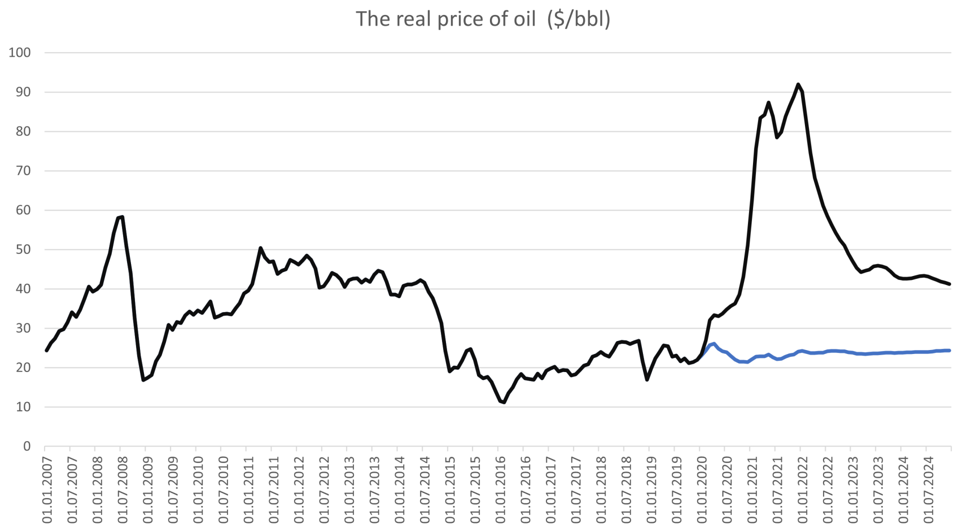

Figure 7: The real price of oil in $/bbl. Forecasted price of oil in the zero-restricted scenario (black line). Forecasted price of oil in the mean-restricted scenario (blue line). CPI Deflated.

To examine the impact on the real oil price, we have transformed the logged data back to its original form. The findings indicate that, after adjusting for inflation, the oil price rises from $22/bbl to approximately $90/bbl within a two-year period. This price escalation can be attributed to the shutdown of Norwegian production as well as the overall decline in global production in the zero-restricted scenario. One caveat to this result is that we have imposed a remarkably large shock to a linear model. If the underlying economic processes are nonlinear in the space that we push the model into, then our model's predictions become less accurate. This result should therefore be interpreted as an upper bound of the implied price response. If we look at the 10th percentile forecast for the oil price represented by the orange line in Figure 4, it suggests that the two-year ahead price lies around $48/bbl. This is more in line with reasonable expectations. The price starts converging back after two years and lies around \$40/bbl for the median forecast and at \$30/bbl for the 10th percentile forecast at the end of our horizon. These outcomes are consistent with those observed in the temporary scenarios, but with greater magnitude and durability.

The forecast shows that Real Economic Activity is changing a lot during our horizon. The contemporaneous increase in economic activity and oil price is a bit counter intuitive as one would think that demand falls when oil prices rise. However, economic agents are slow to adjust their behaviour, which can explain the later fall in Real Economic Activity. The median forecast in Figure 5 shows that Real Economic Activity is constant when the oil price is unchanged.

5 Conclusion

As the urgent need to transition away from fossil fuels grows, this thesis has examined Norway's ability to influence the world's oil market. To accomplish this, we first analysed the 1986 oil strike to investigate the response of the global oil market to a decrease in Norwegian oil production during that period. Subsequently, a structural VAR model is utilized to examine the persistence of the findings derived from the 1986 strike within a temporary context. Finally, we employed a conditional forecast to analyse the effect in a permanent shock scenario.

Our results suggest that large temporary shocks to Norwegian production influence the price of oil. The 1986 strike contributed to a 27% rise in the oil price while our structural VAR model reproducing the same shock suggest a 6.5% increase in the price after three months. Regarding global oil production, the temporary shock in the SVAR model gave a short-term accumulated increase in production, but reported no significant change in the long-term. In the case of a permanent shutdown, we saw no effect on global production in the short-term but a positive impact after around nine months. Global production stabilizes at a lower level, leading to beneficial effects on CO_2 emissions from oil. The shutdown also highly influences the price, increasing it from 22$/bbl to 90$/bbl during the first two years, before decreasing to approximately $40/bbl at the end of our horizon. The price if Norway continues production is unchanged.

To take this analysis further, we propose implementing the conditional forecast methodology to a downscaling of Norwegian production instead of a complete shutdown. By doing so, the analysis will better capture the real-world implications and potentially sidestep the issues of nonlinearity. Additionally, we recommend considering the economic impact on Norway before reaching any conclusive decisions regarding the course of action to be taken.

References

Ahlvik, L., Andersen, J. J., Hamang, J. H., & Harding, T. (2022). Quantifying supply-side climate policies. https://biopen.bi.no/bi-xmlui/bitstream/handle/11250/2983258/CAMP_WP_1_2022.pdf?sequence=1&isAllowed=y

Baumeister, C., & Guérin, P. (2021). A comparison of monthly global indicators for forecasting growth. *International Journal of Forecasting, 37*(3), 1276-1295. https://www.sciencedirect.com/science/article/pii/S0169207021000492

Baumeister, C., & Hamilton, J. D. (2019). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review, 109(5), 1873-1910.* DOI: 10.1257/aer.20151569. https://pubs.aeaweb.org/doi/pdfplus/10.1257/aer.20151569

Blake, A. P., & Mumtaz, H. (2017). Applied Bayesian econometrics for central bankers. *Technical Books.* https://www.bankofengland.co.uk/-/media/boe/files/ccbs/resources/applied-bayesian-econometrics-for-central-bankers-updated-2017.pdf

Caldara, D., Cavallo, M., & Iacoviello, M. (2019). Oil price elasticities and oil price fluctuations. *Journal of Monetary Economics, 103, 1-20.* https://www.matteoiacoviello.com/research_files/CCI.pdf

CFR. (2005). Council on Foreign Relations. *Katrina and Oil Prices* https://www.cfr.org/interview/katrina-and-oil-prices#:~:text=Hurricane%20Katrina%20caused%20severe%20damage,Strategic%20Petroleum%20Reserve%20

Doan, T., Litterman, R., & Sims, C. (1984). Forecasting and conditional projection using realistic prior distributions. *Econometric reviews, 3*(1), 1-100.

European Council. (2022). *Council agrees on the Carbon Border Adjustment Mechanism (CBAM).* https://www.consilium.europa.e u/en/press/press-releases/2022/03/15/carbon-border-adjustment-mechanism-cbam-council-agrees-its-negotiating-mandate/

Federal Reserve. (2013). *Oil Shock of 1978–79* https://www.federalreservehistory.org/essays/oil-shock-of-1978-79

Fæhn, T., Hagem, C., Lindholt, L., Mæland, S., & Rosendahl, K. E. (2017). Climate policies in a fossil fuel producing country–demand versus supply side policies. *The Energy Journal, 38*(1). https://www.nmbu.no/sites/default/files/ej_2017_faehn_hagem_lindholt_maeland_rosendahl_demand_vs_supply_side_climate_policies_norway.pdf

Hamilton, J.D. and Herrera, A. (2004) Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy. *Journal of Money, Credit, and Banking, 36*, 265-286.https://doi.org/10.1353/mcb.2004.0012

Hamilton, J. D. (1983). Oil and the macroeconomy since World War II. *Journal of political economy, 91*(2), 228-248. https://www.journals.uchicago.edu/doi/epdf/10.1086/261140

Hamilton, J. D. (1985). Historical causes of postwar oil shocks and recessions. *The Energy Journal, 6*(1). 97-116. https://www.jstor.org/stable/pdf/41322100.pdf?refreqid=excelsior%3Ab7d0295439fd127ef83ce2b2c3ce07dc&ab_segments=&origin=&initiator=

IEA. (2021). *Greenhouse Gas Emissions from Energy Data Explorer.* https://www.iea.org/data-and-statistics/data-tools/greenhouse-gas-emissions-from-energy-data-explorer

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review, 99*(3), 1053-69. ISSN: 0002-8282. DOI: 10.1257/aer.99.3.1053. https://www.aeaweb.org/articles?id=10.1257/aer.99.3.1053

Liston, R. (1986, 25. April). Norway oil strike ends; oil prices plummet. *United Press International.* https://www.upi.com/amp/Archives/1986/04/25/Norway-oil-strike-ends-oil-prices-plummet/1631514789200/

Lohr, S. (1986, 7. April). Strike Idles Norwegian Oil Production. *New York Times.* https://www.nytimes.com/1986/04/07/business/strike-idles-norway-oil-production.html

Lyngve, E. (1986). Strike shuts down oil industry. *United Press International.* https://www.upi.com/Archives/1986/04/06/Strike-shuts-down-oil-industry/4355513147600/

Malkin, C. (2020). *Gibbs Sampling.* https://towardsdatascience.com/gibbs-sampling-8e4844560ae5

Meland, T. (2020). Streik og lockout. *Industriminne Statfjord.* https://statfjord.industriminne.no/nb/2019/12/03/streik-og-lockout/

Nasjonalbiblioteket. (2023). *Nettbiblioteket* https://www.nb.no/search?q=olje%20streik&mediatype=aviser&fromDate=19860301&toDate=19860408

Norges bank. (2019). *Dette er Oljefondet.* https://www.nbim.no/no/oljefondet/om-oljefondet/

Norsk Petroleum. (2023). *Alt du trenger å vite om norsk petroleumsvirksomhet.* https://www.norskpetroleum.no/

NRK. (2021). *Olje.* https://www.nrk.no/valg/2021/partiguiden/nb/tema/olje/

Rebei, N., & Sbia, R. (2021). Transitory and permanent shocks in the global market for crude oil. *Journal of Applied Econometrics, 36*(7), 1047-1064. https://onlinelibrary.wiley.com/doi/epdf/10.1002/jae.2863

Retriver. (2023). *Atekst.* https://app.retriever-info.com/services/archive?canFetchDataOnDateSelectorChange=true&fromDat

e=1986-03-01&searchString=olje%20streik&toDate=1986-04-08

Riekeles, H. (2023). Norsk olje, globale utslipp. https://www.vista-analyse.no/site/assets/files/7858/va-rapport_2023-04_norsk_olje_-_globale_utslipp.pdf

Ryggvik, H. & Rosendahl, K. E. (2021). Nedtrappinge av norsk olje og gass. *Grønn Industri 21.* http://manifesttankesmie.no/wp-content/uploads/sites/2/2021/05/Rapport-5-Nedtrappingsstrategi-av-norsk-olje-og-gass.pdf

Rystad Energy. (2023). Netto klimagassutslipp fra økt olje- og gassproduksjon på norsk sokkel. https://www.regjeringen.no/contentassets/f5fc522f50674c1f9e0b5db47c264dbe/netto-klimagassutslipp-fra-okt-olje-og-gassproduksjon-pa-norsk-sokkel_hovedrapport.pdf

Sims, Christopher A. (1980). 'Macroeconomics and Reality.' *Econometrica*, 48.1, pp. 1–48. ISSN: 00129682, 14680262 http://www.jstor.org/stable/1912017.

United Nations. (n.d). *Causes and Effects of Climate Change.* https://www.un.org/en/climatechange/science/causes-effects-climate-change

Waggoner, D. F., & Zha, T. (1999). Conditional forecasts in dynamic multivariate models. *Review of Economics and Statistics, 81*(4), 639-651. https://www.jstor.org/stable/pdf/2646713.pdf?refreqid=excelsior%3A9c4d1677f821328be0f3f045304c1dd1&ab_segments=&origin=&initiator=

Woldsdal, M. (1986, 25. March). Ministermøte avsluttet uten resultater: Prisfall efter OPEC fiasco. *Aftenposten Morgenblad.* https://app.retriever-info.com/services/archive?canFetchDataOnDateSelectorChange=true&fromDate=1986-03-25&searchString=OPEC&toDate=1986-03-26

A Appendix

A.1 Section 3 - Plot of the WTI futures market 1986

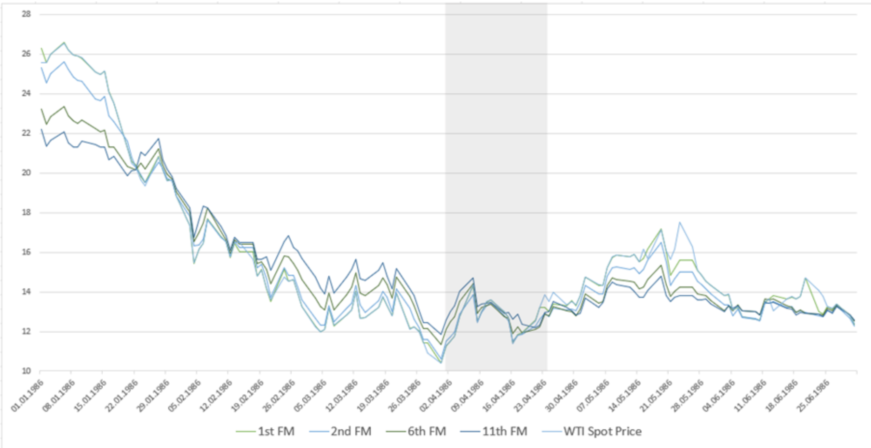

Figure 8: Plot of the 1st, 2nd, 6th, and 11th position of the NYMEX WTI futures prices in deviations from the WTI spot price for Crude oil. 01.01.86–25.06.08, daily frequency. The shaded area marks the 1986 strike. Y-axis: deviation from spot in dollars. Data is retrieved from Macrobond.

A.2 Section 4.1.2 - The Data

Variable	Description
$\Delta nor.prod$	Norwegian crude oil production, including lease condensate, in thousands of barrels per day. Retrieved from EIA. (We use log-differenced data in the SVAR model, and raw data for the conditional forecast.)
$\Delta g.prod$	Global crude oil production, including lease condensate, in thousands of barrels per day. Retrieved from EIA. Norwegian production has been subtracted from global production. (We use log-differenced data in the SVAR model, and raw data for the conditional forecast.)
rea	Monthly World Industrial Production (WIP) index. Source: Baumeister, C. and J.D. Hamilton (2019). Data available on Baumeister's personal website: `https://sites.google.com/site/cjsbaumeister/datasets` We chose to use the 24-month growth of the WIP index as it provides a longer-term perspective on global activity. This approach smoothes out short-term fluctuations and provides a more stable measure of overall industrial production trends, which is useful when focusing on long-term patterns.
$lrpo$	CPI deflated U.S. crude oil imported acquisition cost by refiners (the real price of oil), logged. Retrieved from EIA.

Table 1: The data

The data is in monthly frequency and goes from 1974:M01–2019:M12 for both the structural VAR and the conditional forecast. We chose to exclude data from 2020–2023 to avoid the exceptional disruptions caused by the COVID-19 pandemic, allowing for a clearer analysis of normal market conditions and long-term trends.

We have taken the first difference in oil production data to transform it into a stationary series to ensure the invertibility of the reduced form coefficient matrix and promote the stability of the VAR system. This transformation helps eliminate non-stationarity, enabling more reliable and robust analysis of the relationships between variables in the model. The data is logged as it allows for interpretation in terms of percentage changes, stabilizes variances, linearizes nonlinear relationships, and facilitates the analysis of elasticities.

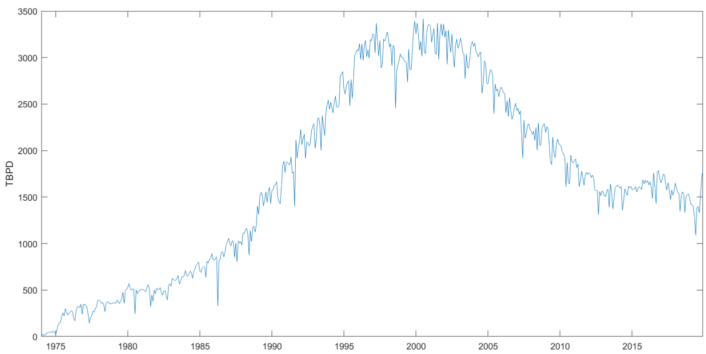

Figure 9: Plot of Norwegian crude oil production, including lease condensate. 1974:M01–2019:M12. Source: EIA.

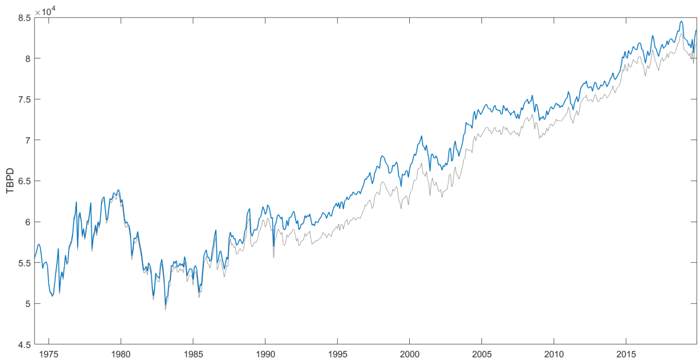

Figure 10: Plot of Global crude oil production, including lease condensate (Blue line). Plot of Global Production - Norwegian Production (Grey line). 1974:M01–2019:M12. Source: EIA.

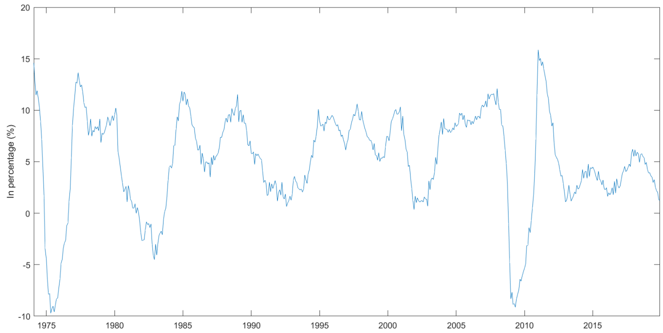

Figure 11: Monthly World Industrial Production (WIP) index. 1974:M01–2019:M12. Source: Baumeister, C. and J.D. Hamilton (2019).

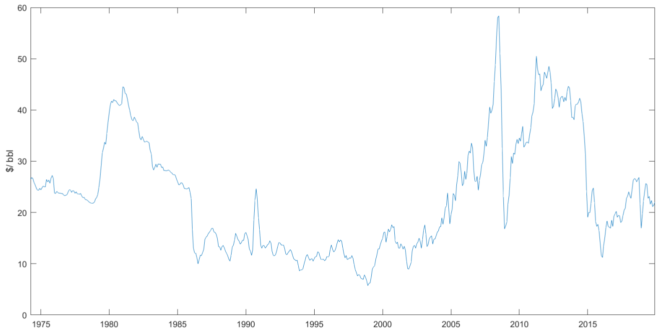

Figure 12: Plot of the CPI deflated U.S. crude oil imported acquisition cost by refiners (the real price of oil), 1974:M01–2019:M12. Source: EIA.

A.3 Section 4.1.3 - Impulse responses from structural VAR

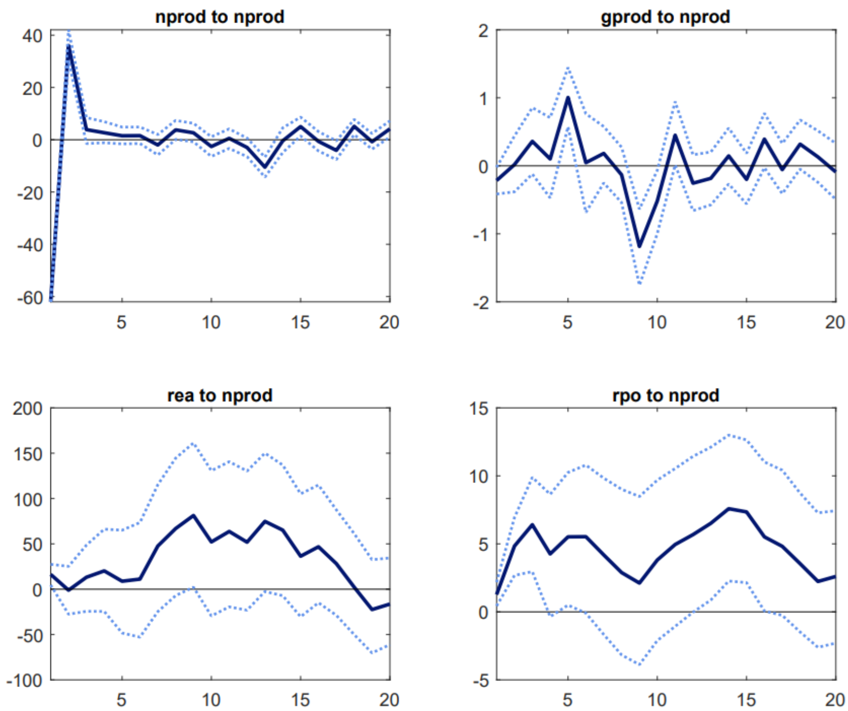

Figure 13: Impulse responses of a 62% temporary reduction in Norwegian oil production. Log-differenced production data, not accumulated. Confidence Level: 68%.

Following the 62% shock, production recovers quickly and surpasses its original level in the second period. After three periods, Norwegian production returns back to normal. This is slightly different from the 1986 strike, where production did not exceed its original level after the shock.

A.4 Section 4.2.1 - Unconditional forecast computation

$\tilde{X}_{t+1}, \tilde{X}_{t+2}, \ldots, \tilde{X}_{t+k}$ are computed by first making an $(horizon + L \times N)$ matrix of zeros. We use 24 lags also for the conditional forecast, making the matrix (84×4). Next, we initiated a loop to generate forecasts for each period beyond the observed data. Within this loop, lagged values of the forecasted endogenous variables are stored in a matrix. To capture the lagged relationship between the variables, the loop iterates over the lags for each period and appends the lagged values of the forecasted variables to the matrix. To construct the forecast for the current period, the stacked matrix is multiplied by the coefficient matrix, which contains the estimated coefficients of the VAR model. This multiplication yields the forecasted values for the current period. The forecasted values are then assigned to the corresponding row of the first generated matrix, this is repeated for each period in the forecast horizon. After the loop is complete, the initial lagged periods, which were initialized with observed data, are removed from the forecast. This step ensures that only the unconditional forecast values remain.

The resulting matrix contains the unconditional forecast values for all four variables in Y_{t+K}. In our r matrix, we have used the values for Norwegian production, which corresponds to the first column of the matrix.

A.5 Section 4.2.1 - Gibbs sampling algorithm

Step 0) Create parameters.

First, create a parameter that determines the total number of iterations for the algorithm ($REPS$) and a parameter that determines the number of initial iterations to discard as the burn-in period ($BURN$). Second, create empty matrices to hold the forecasts of the variables ($out1, \ldots, out4$). Third, initialize the conditional forecasts ($\hat{Y}_g$) and the error variance (σ).

Step 1) Initialise the VAR coefficients and the impact matrix (A_0).

First, extend the dataset to include the forecasted values ($Y_t^\star = [Y_t; \hat{Y}_{t+k}]$). By doing this, the draws of the VAR parameters consider the restrictions imposed on the forecasts ($R\varepsilon = r$). Second, take lags of the data to construct the input matrix ($X_t^\star$). Then, calculate the conditional mean (M) and variance (V) by implying $M = \vec{X}^\star / \vec{Y}^\star$ and $V = \sigma \otimes (X^{\star\prime} X^\star)^{-1}$.[17] After that, draw the VAR parameters from a multivariate normal distribution. Lastly, draw σ from the inverse Wishart (IW) distribution, and calculate A_0 by performing Cholesky on σ. Waggoner and Zha (1999) show that the choice of identification strategy does not affect the conditional forecast as it only depends on the reduced VAR. Cholesky is therefore chosen for calculating A_0 out of simplicity.

Step 2) Construct unconditional forecast, the R matrix, and the r matrix.

First, compute the impulse responses by using the updated impact matrix and the VAR parameters drawn in *step 2*. Note that the matrix

[17] The Kronecker product, denoted by $\otimes$, is a binary operation that takes two matrices and produces a larger matrix by combining each element of the first matrix with the second matrix.

storing the parameters has to be reshaped to a $(NL + 1 \times N)$ matrix. Second, calculate the unconditional forecast in the same way as shown in Appendix A.4 by using the new VAR parameters. Then, construct a new R matrix using the new responses and a new r matrix using the new unconditional forecast. Now, compute the mean and variance of the distribution of the restricted structural shocks (See, Equation 8).

Step 3) Construct conditional forecast.

First, draw the structural shocks from the $N(\bar{M}, \bar{V})$ distribution and reshape the draws to contain $(N \times Horizon)$ elements. Then, compute the conditional forecast by using the same approach as for the unconditional forecast, but include the drawn restrictions in combination with the new impact matrix.

The forecasted values for each variable are then appended to their respective outcome matrices. The algorithm runs for an $REPS$ amount of iterations and updates the error covariance and VAR parameters. For each time, new forecasts are generated after the burn-in stage $BURN$. The final forecasts are stored in the outcome $(out1, \ldots, out4)$ matrices after the algorithm is complete.

A.6 Section 4.2.1 - Posterior distribution

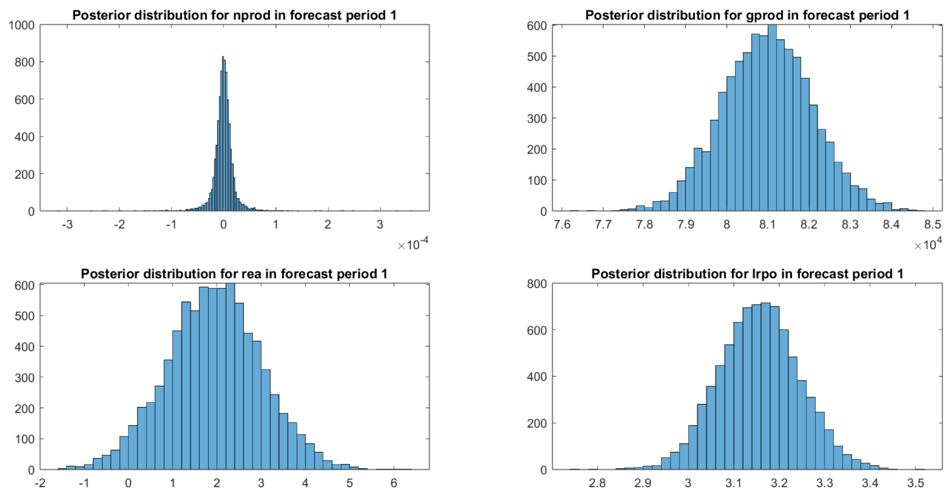

Figure 14: Posterior distribution for the first forecast period in the zero-restricted scenario.

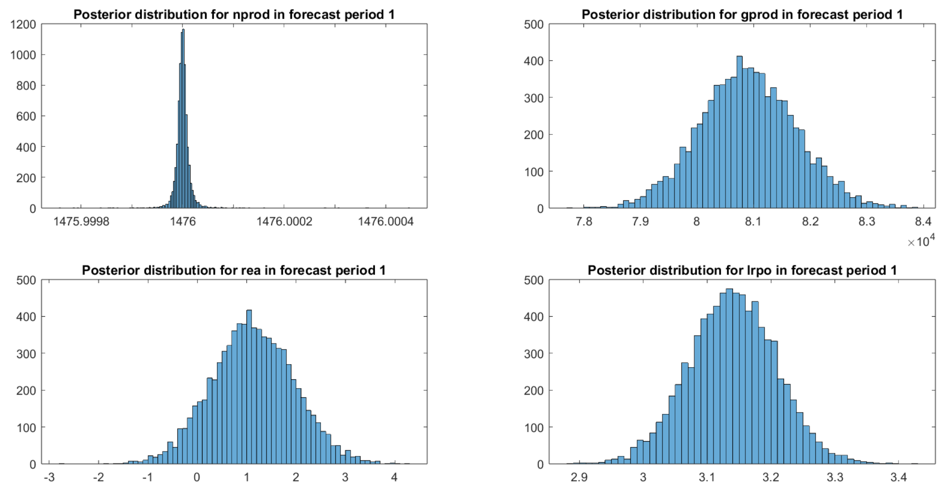

Figure 15: Posterior distribution for the first forecast period in the mean-restricted scenario.

A.7 Section 4.2.2 - Plot of results (CF)

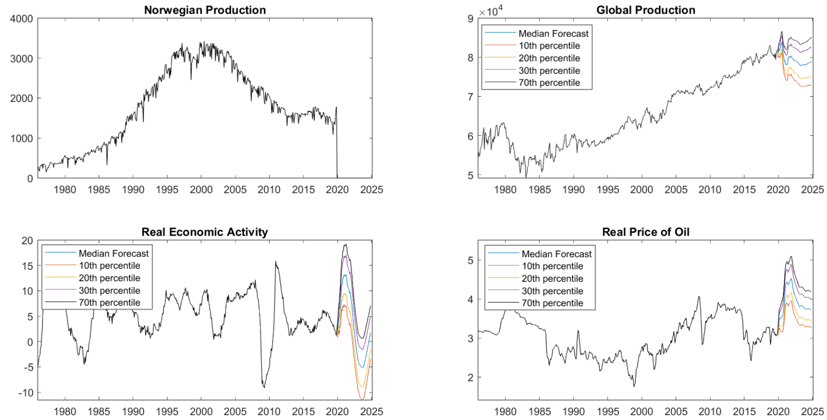

Figure 16: Plot of CF in the zero-restricted scenario. Dataset: 1974:M2–2019:M12. Forecast horizon: 60 periods (five years).

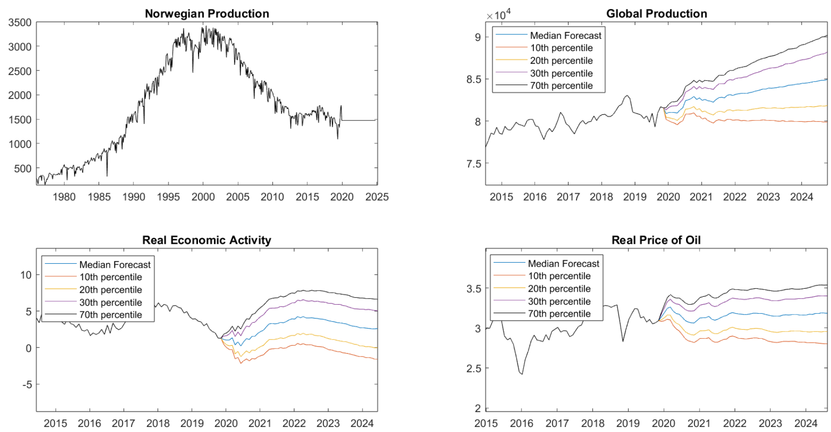

Figure 17: Plot of CF in the mean-restricted scenario. Dataset: 1974:M2–2019:M12. Forecast horizon: 60 periods (five years).

List of Figures

List of Tables