

DIPLOMARBEIT

What Drives Drilling Up and Prices Down?

A Structural Vector Autoregressive Model of the U.S. Natural Gas Market

zur Erlangung des akademischen Grades

Diplom-Ingenieur/in

im Rahmen des Studiums

Statistik-Wirtschaftsmathematik

eingereicht von

Valentin Elias Winkler, BSc Matrikelnummer 11811850

ausgeführt am Institut für Stochastik und Wirtschaftsmathematik der Fakultät für Mathematik und Geoinformation der Technischen Universität Wien

Betreuung

Betreuer: Associate Prof. Dipl-Vw. Emanuel Gasteiger PhD

Wien, 23. Mai 2023

Abstract

This study uses a structural vector autoregressive (SVAR) model to examine the relationships between the intensity of drilling activity (i.e. investment) for natural gas production, natural gas withdrawals, economic activity, and natural gas prices in the United States. The results show that the reaction of drilling to an unexpected change in natural gas prices depends on the source of the price change. Specifically, I find that the reaction of drilling is significantly stronger after an economic activity shock than after a gas demand shock (e.g. due to oil price fluctuations). In addition, it is shown that demand-side factors were more important than supply-side factors in explaining the 85% drop in natural gas prices from June 2008 to April 2012. This contradicts prevailing explanations focused on shale gas development and should dampen the expectations of policymakers seeking to rapidly expand shale gas production in order to obtain similar cheap energy as the U.S. after 2008.

Acknowledgements

There are many people without whom working on this thesis and conducting my studies would have been way harder and less enjoyable than has been the case. First of all, I am extremely grateful to Emanuel Gasteiger. He is not only the supervisor of this thesis but has supported me professionally, intellectually, and personally on countless occasions over the last few years. To give some examples directly related to this work, Emanuel was the first to introduce me to the topic of structural vector autoregressions, suggested I apply for the internship during which most of the work that went into this thesis was done, and later organized a seminar presentation of my project. Through it all, he was always down-to-earth, thoughtful, and understanding.

Another person I am deeply indebted to is Sven Schreiber, who was my adviser during the internship at the IMK in Düsseldorf last summer, where I developed the project that led to my thesis. Although I had hardly any experience in applied time series analysis before, he gave me a lot of time and freedom to familiarize myself with the topic, try things out and find a research question that was interesting for me. At the same time, he always patiently helped me with questions and also took the time to discuss fundamental matters of the subject with me. Even after my internship was over, Sven still gave me detailed feedback, which greatly improved my manuscript. Of course, I alone am responsible for any remaining mistakes and weaknesses.

I am fortunate to have had an excellent environment to study and work on my thesis. The Economics Research Group at TU Vienna, where I have worked as a student assistant for the past two and a half years, has ensured that I have always felt part of the community. I am grateful to my colleagues for their help, for all the interesting discussions, and for making the lunch breaks so enjoyable. Furthermore, I would also like to thank the people at the IMK in Düsseldorf for providing me with such a good environment to work on my project and for giving me an insight into their interesting work. During my studies, I had many great teachers, of whom I cannot list all. However, I would like to explicitly thank Wolfgang Scherrer, whose lectures on time series econometrics gave me an in-depth introduction to the subject and whose accurate and aesthetically pleasing approach to econometrics I have come to appreciate very much.

The fact that the half-decade during which I studied at the TU Vienna was such a nice one is mainly due to my friends and I would like to thank each and everyone for their time and support. I am especially grateful to Jakob, without whom I probably would not have dared to start studying mathematics and who accompanied me through the hardest and most pleasant parts of my studies. Since we both moved to Vienna many things have changed, but our friendship has remained a constant.

The person who has influenced me the most over the last two years is my girlfriend Miriam. I am infinitely grateful for her unconditional support, honest criticism, and enrichment of my ideas, her sharp mind, and her cheerful nature. I would especially like to thank her for the visits to Düsseldorf during my internship there. Fortunately, I can experience the coming adventures together with her.

Finally, I would like to express my gratitude to my family. I am lucky to have two such smart, reliable, and kind brothers as Konstantin and Lorenz. I am thankful to my parents not only for all the bills they paid during my studies and before but also for always reminding me that there may be things more important than my professional and academic achievements. Thanks to them, I know I have a place where I am always welcome no matter what. I understand more and more that I cannot take this for granted.

Contents

1	Introduction						
2	Literature Review						
3	The Shale Gas Boom						
4	The Model4.1Econometric Set Up4.2Choice of Variables4.3Estimation of the Reduced Form4.4Identification of the Structural Form4.5Inference in the Structural Model	 13 13 15 18 20 24 					
5	Results5.1Comparison of Identification Strategies5.2Natural Gas Market Dynamics5.3The Price Elasticity of Natural Gas Drilling5.4Forecast Error Variance Decomposition5.5Historical Decomposition in the Presence of Unit Roots5.6Decomposition of Drilling Activity5.7Disentangling Historical Natural Gas Price Fluctuations	26 28 30 31 35 36 37					
6	Conclusion 4						
7	List of Figures 4						
8	Bibliography 4						
\mathbf{A}	Mathematical Appendix						
в	3 Replication of the Kilian and Inoue (2013) Oil Model						
С	Robustness Checks C.1 Results for Subsamples C.2 Different Elasticity Bounds	58 58 61					

1 Introduction

After decades of decline and stagnation, starting around 2005 natural gas output in the United States began to rise dramatically. Mainly responsible for this was the exploitation of unconventional natural gas reservoirs, commonly referred to as shale gas. As Mu [2019] and Wang and Krupnick [2015] point out, the process that led to this sudden surge in gas withdrawals actually started in the 1970s and 1980s but came to fruition some decades after. Its impressive success in raising domestic natural gas production also aroused interest of other governments - for example in Argentinia, Poland or Australia - to replicate the U.S. shale gas boom [Wang and Krupnick, 2015]. Responsible for this is the notion that the shale gas boom brought cheap, domestically produced energy to industrial and residential users. Starting with the financial crisis of 2008, U.S. natural gas prices collapsed and never rose to the levels experienced in the 2000s: In June 2008, one million British termal units (from now on abbreviated MMBtu, one of which equals roughly 0.29 MWh) of natural gas cost 12.7 USD at the Henry Hub, which is the most important trading platform for this fuel in the United States, while in April 2012 prices were as low as 1.95 USD.¹ After 2008 natural gas did not only become cheaper but also increasingly important in the United States, surpassing coal as the most used fuel for power generation in 2016 [Walton, 2017]. The ever-increasing role of natural gas makes it ever more essential to understand the dynamics of natural gas markets. Studying the example given by the United States is not only of relevance to the U.S. itself but also to other countries, who (like the European Union) recently switched from a pricing policy linked to crude oil to an US-style model of gas on gas competition or consider doing so [IGU, 2021]. This puts some questions at the center of interest: To what extent do supply and demand side factors influence gas prices? And how do changes in natural gas prices influence gas production and investment behavior – in particular drilling activity? The importance of the latter question is highlighted by recent studies of oil [Anderson et al., 2018] and natural gas wells [Mason and Roberts, 2018]. Both papers show that the rate of drilling is the most important decision variable for oil and gas producers, since production from existing wells is primarily governed by laws of geology and petro-engineering. Nevertheless, the issue of oil and gas drilling has been mainly dealt with in a narrow setting focussing on supply and demand conditions within energy markets, leaving out general macroeconomic conditions, while macroeconomic gas market models usually leave out drilling as a variable.

I try to close this gap by setting up a structural vector autoregressive (SVAR) model containing drilling activity, gas production, industrial production and the Henry Hub gas price. This model is able to verify previous findings of related studies like Arora and Lieskovsky [2014], Wiggins and Etienne [2017] and Rubaszek et al. [2021], who have already shown that demand side shocks are way more important than supply side shocks for explaining variation in the natural gas price. Moreover, according to my model drilling shocks have no significant effect on other gas market specific variables. Most likely this is because drilled wells usually last decades, often times longer than the sample used in this study, and therefore the effect on production is streched out over time and hard to capture for a monthly SVAR. Drilling activity, on the other hand, seems to rise when prices rise. This response, however, takes some time and depends on the origin of the price shock. The price elasticity of drilling after an economic

¹Source: FRED (2022) 'Henry Hub Natural Gas Spot Price, Not Seasonally Adjusted (MHHNGSP)' (monthly): https://fred.stlouisfed.org/series/MHHNGSP.

activity shock is significantly higher than after a gas demand shock.² The reason for this is that the effect of a gas demand shock on natural gas prices vanishes quickly, while the effect of an economic activity shock is persistent. In particular I show that the extraordinarily high level of active rigs before the financial crisis of 2008 was to a large extent fuelled by economic activity shocks. After this, advances in drilling technology prevented the number of active rigs to collapse despite exceptionally low natural gas prices.

The price decline of U.S. natural gas after its height in 2008 is of particular interest to policymakers in the U.S. and in other countries that consider shale gas development. In April 2012 natural gas prices in the U.S. were around 85 percent below their June 2008 level. This drop stood in sharp contrast to rising global LNG prices due to increased demand after the Fukushima accident and extremely high crude oil prices. The United States seemed to have found a recipe to escape the energy price bonanza. It is tempting to attribute the extremely low natural gas prices predominantly to the shale revolution. Wiggins and Etienne [2017] and Rubaszek et al. [2021] have already shown that demand side factors are important to explain the price drop after 2008. However, they do not calculate the relative contributions of different shocks to the change in the price level. I estimate that between June 2008 and September 2009 more than two thirds of the decline in natural gas prices have to be attributed to gas demand and economic activity shocks. Perhaps more surprisingly, when comparing April 2012 prices to June 2008 prices, demand side factors were still more important than supply side factors. One should not overstate the effect of the shale boom on natural gas prices in that period and instead take into account low energy demand due to the financial crisis. Of course, the fact that the rise in shale gas production made the U.S. virtually independent of global LNG markets might have increased its resilience against price shocks on the world market. Most likely, the combination of weak domestic demand and global price shocks to which the U.S. was relatively immune led to the striking difference between natural gas prices in the United States and in Europe as well as Asia that could be observed after the financial crisis.

In the following sections I will give a brief review of the relevant literature and provide some descriptive evidence of U.S. natural gas markets in the shale boom era. After this I will set up a SVAR model of the U.S. natural gas market that is identified by a combination of sign and exclusion restrictions. Inference will be conducted using a modified version of the Bayesian approach developed by Inoue and Kilian [2013, 2018] for purely sign identified models. As a next step I will present estimates for the impulse response functions, decompose the forecast error variance of the endogenous variables and derive a historical decomposition of natural gas prices and drilling activity between 1993 and 2020. The final section concludes. Mathematical derivations of claims made in the paper and some sensitivity analysis are given in the appendix.

2 Literature Review

Oil market literature. The literature on natural gas markets draws heavily on methods and discussions previously developed with respect to oil markets. Since Kilian [2009] SVARs are the predominant tool for analyzing the latter. Using recursive exclusion restrictions and modelling oil production, economic activity and oil prices, this influential study came to the conclusion that demand for oil stemming from changes in global economic activity was mainly responsible

²I define this elasticity as $\partial \log(rig_{t+h})/\partial \log(p_t)$, where rig_{t+h} is the number of active rigs in month t + h and p_t is the real henry hub price in month t.

for rising oil prices in the 2000s, which explained why this price increase seemed to not curb economic growth. Before, studies like Hamilton [2003] focused mainly on oil price shocks arising due to supply disruptions. Kilian and Murphy [2012] later developed the basic model further so that it could be identified using sign and elasticity constraints instead of exclusion constraints, while Kilian and Murphy [2014] added oil inventories to the model to capture speculative demand. Recently Baumeister and Hamilton [2019] have argued that the lower bound for the price elasticity of oil supply used by Kilian and Murphy [2014] is too close to zero. Imposing an explicit prior distribution on this value that allows for larger posterior estimates, using a different measure of global economic activity and accounting for measurement errors of oil inventories, they estimate the importance of supply shocks for explaining oil price fluctuations to be higher than in the model of Kilian and Murphy. The whole methodological debate about how to correctly identify structural shocks in the oil market is too vast to be summarized here. Various strands of this debate address, for example, what prior assumptions can be made about the price elasticity of oil supply, how to measure global economic activity, and how to define particular elasticities. Kilian and Zhou [2020] give an overview of these discussions.

Gas market models. In the last decade there has also evolved a literature on gas market SVAR models. Nick and Thoenes [2014] set up such a model for the German market using natural gas, oil and coal prices, LNG imports to Europe, gas storage levels and a supply shortfall variable. Since their specification uses weekly data they are not able to include any variables concerning real economic activity. They find out that oil and coal prices are an important determinant of natural gas prices in the long run. Arora and Lieskovsky [2014] apply the specification of Kilian [2009] to the U.S. natural gas market and include residential natural gas consumption as a fourth variable. They report evidence for a magnified response of industrial production to a gas supply shock in the shale gas period. Moreover, according to their analysis, supply and business cycle shocks play only a marginal role in explaining natural gas price fluctuations. Hou and Nguyen [2018] detect four regimes in the U.S. gas market and set up a Markov switching SVAR model. They find out that gas demand shocks are the most important determinants of the gas price while gas supply is insensitive to demand side shocks. Rubaszek and Uddin [2020] use a threshold SVAR model to estimate effects of structural shocks depending on inventory levels. All three studies previously mentioned rely on a recursive identification strategy.

Wiggins and Etienne [2017] use the sign restrictions of Kilian and Murphy [2014] combined with a time varying parameter (TVP) and stochastic volatility specification. They report somehow higher magnitudes of the price elasticities of supply and demand at the end of their sample when compared to the early 2000s. Furthermore, they conclude that demand shocks are important to explain the collapse of natural gas prices after 2008. While the flexibility of TVP models is appealing, there is also some potential for criticism. Kilian and Zhou [2020] highlight that even in the absence of structural change a TVP VAR would report time varying parameter values since the degree of over-fitting is governed by hyperparameters for which there is no empirically observable value. Having this in mind, the parameters reported by Wiggins and Etienne [2017] seem to be stable enough during the shale gas boom to not invalidate the use of constant parameter models. Finally, Rubaszek et al. [2021] apply the oil market model of Baumeister and Hamilton [2019] to the gas market of the United States. This approach relies on imposing prior assumption directly on the structural parameter matrix A_0 , the values of which have an economic interpretation as contemporaneous elasticities. This is motivated by the result of Baumeister and Hamilton [2015] that by imposing a flat prior on the orthonormal matrix Q, as is common for sign restricted models, and calculating A_0^{-1} as $chol(\Sigma)Q$, where Σ is the variance-covariance matrix of the reduced form residuals, one imposes an informative prior on A_0 . Therefore this common prior used for sign restricted models is not agnostic, as is often intended. Baumeister and Hamilton therefore suggest to be explicit about the prior one imposes on A_0 . However, as [Kilian and Lütkepohl, 2018, 453-454] highlight, a flat prior on A_0 can lead to an unintentionally informative prior on A_0^{-1} . Therefore choosing between the two approaches should depend on which variable is of interest. If this are the elasticities given by A_0 , the Baumeister and Hamilton [2019] approach might be preferable, while when needing the values of A_0^{-1} , as is the case for computing the structural impulse response functions, the forecast error variance decomposition or the historical decomposition, the 'traditional' approach might be better suited. Since I am mainly interested in A_0^{-1} , I choose the latter setting.

Drilling literature. Against the backdrop of major improvements in (horizontal) oil and gas drilling in the last two decades, there has been a growing empirical and theoretical literature on oil and gas drilling. There are three major contributions of Ryan Kellogg, one of which is with coauthors: Kellogg [2011] shows how inter-firm learning increases the productivity of certain matches between fossil fuel companies and drilling contractors, which can explain the stability of contract relationships and the consistent improvement of productivity in oil and natural gas drilling. Kellogg [2014] uses data on oil drilling in Texas to show how oil price volatility influences drilling negatively, the response being consistent with optimal behaviour according to a theoretical model he proposes. Moreover, volatility measures derived from futures prices are a better predictor for drilling than variation derived from past prices, for example by means of fitting a GARCH process. Finally, Anderson et al. [2018] show that oil production from existing wells in Texas is unresponsive to oil prices, which can be explained by laws of geology and petro-engineering. Drilling, on the other hand responds strongly to oil price changes. A similar result for natural gas is provided by Mason and Roberts [2018], who study natural gas wells in Wyoming and report price elasticities of supply that are small in magnitude. The last two papers are crucial for my work, since they motivate including drilling activity into the gas market model and justify some of my identifying assumptions. Further evidence on the relationship between prices and drilling is provided by Khalifa et al. [2017] as well as Shakya et al. [2022]. Using quantile regression techniques, the former show that there is a strong relationship between the rig count and oil prices in the US, especially between lagged prices and the rig count. The authors conclude that the relationship between rig count and oil prices is not contemporaneous but that the reaction of drilling to price changes takes time. The latter article conducts Granger causality analysis to study the relationship between oil drilling, gas drilling, oil prices and gas prices. The authors use a rolling window approach to account for structural change and find that information transmitted between drilling and prices has increased since the development of shale oil and gas. They also note a growing importance of gas drilling and prices as information transmitters.

All in all, there seems to be a lack of structural macro-econometric models trying to explain the dynamics governing and effects of oil or gas drilling. Existing papers on drilling mainly focus on the micro-level or conduct traditional time series analysis using reduced form equations and ignoring general macroeconomic conditions or treating them as exogenous. Structural gas or oil market models, on the other hand, usually do not consider drilling. This paper tries to close some part of this gap.



Figure 1: Natural gas withdrawals in the United States and the contribution of shale gas. Source: Energy Information Administration (2022) 'Natural Gas Gross Withdrawals and Production' (annual): https://www.eia.gov/dnav/ng/ng_prod_sum_dc_NUS_mmcf_a.htm.

Notes: As is common in the United Sates, natural gas withdrawals are denoted in cubic feet (CF), one of which roughly equals 0.028 cubic meters.

3 The Shale Gas Boom

Until less than two decades ago, natural gas output in the USA had been stagnant for a long time. This changed dramatically with the shale gas revolution. Shale gas is natural gas contained in shale – a sedimentary rock with low permeability. Because of the latter property, shale gas does not naturally flow into a drilled well but needs additional stimulation by hydraulic fracturing ('fracking') [Mu, 2019]. Even though shale gas production in the U.S. ranges back to the 19th century and gained political support after the oil crisis of the 1970s, it needed some more decades and advances in horizontal drilling as well as hydraulic fracturing to become economically viable. According to Wang and Krupnick [2015] it was a combination of government policy, high natural gas prices in the 2000s, certain institutional characteristics of the U.S. oil and gas industry, environmental factors, deep capital markets and a dynamic oil and gas service industry that finally led to the shale revolution in the 2000s. As Figure 1 shows, the term 'revolution' is appropriate when talking about the roaring success of shale gas. In 2007 the share of natural gas withdrawals from shale wells was rather negligible, while some years after more gas was withdrawn from shale wells than from conventional wells. This rise in shale production also coincided with a rise in total U.S. natural gas production. As is shown in Figure 2, this facilitated a changing role of the United States on the international natural gas markets. Before 2007 net natural gas exports from the U.S. had been declining, but after 2007 this trend changed. While imports were reduced, exports (mainly to Mexico) rose, pushing net exports upwards. After the opening of the natural gas liquefaction plant Sabine Pass I in 2016 exports rose even further to finally surpass total imports. Since then, the United States have been a net exporter of natural gas. Figure 3 sheds some light on the liquefied natural gas (LNG) trade balance of the United States. In the 2000s the U.S. started importing substantial amounts of LNG. After the start of the shale gas revolution and the great recession of 2008, LNG imports declined steeply. When the U.S. started to open their own liquefaction facilities in 2016, which they have expanded rapidly since then, they quickly became a net exporter of LNG. It is notable, that apart from the Covid crash the liquefaction plants were always close



Figure 2: Natural gas production in the United States and international gas trade. Source: Energy Information Administration (2022) 'Natural Gas Gross Withdrawals and Production' (monthly); 'U.S. imports by country' (monthly); 'U.S. exports by country' (monthly): https://www.eia.gov/naturalgas/ data.php.

Notes: Lines mark takeoff of U.S. shale gas production (2007) and opening of liquefaction plant Sabine Pass I (2016).



Figure 3: LNG trade and liquefaction capacity in the United States. Source: Energy Information Administration (2022) 'U.S. liquefaction capacity': https://www.eia.gov/naturalgas/data.php.

Notes: Lines mark takeoff of U.S. shale gas production (2007) and opening of liquefaction plant Sabine Pass I (2016).

to full capacity.

During the shale gas boom, not only natural gas production but also nominal and relative prices of natural gas compared to crude oil changed. As is apparent from Figure 4 and noted by Brown and Yücel [2008], during the early 2000s a 1:6 rule did properly approximate the gas price in terms of the oil price. This makes sense, since 6 MMBtu of natural gas contain roughly the same energy as one barrel of crude oil. Shortly before the shale gas take-off, however, this rule seemed to break down. A straightforward explanation for this phenomenon is that additional shale oil production can easily be sold abroad while U.S. natural gas exports to



Figure 4: Comparison of crude oil prices (WTI) and natural gas prices (Henry Hub). Source: Energy Information Administration (2022) 'Natural Gas Spot and Futures Prices (NYMEX)' (monthly): https://www.eia.gov/dnav/ng/ng_pri_fut_s1_d.htm; FRED (2022) 'MCOILWTICO' (monthly): https: //fred.stlouisfed.org/series/MCOILWTICO.

Notes: Lines mark takeoff of U.S. shale gas production (2007), oil price collapse due to increased OPEC+ production (2014) and opening of LNG liquefaction plant Sabine Pass I (2016). Crude price is given in USD/Barrel, natural gas price is given in USD/6MMBtu which represents roughly the same energy content [Brown and Yücel, 2008].

other continents are constrained by LNG liquefaction capacity – which did not exist or operate almost at full capacity in the plotted sample. This change in relative prices was also reflected in the pattern of drilling, even though with some considerable lag. Figure 5 plots the number of rigs that are employed for oil or natural gas drilling at a certain point in time. Rigs are tools that can be used to drill new wells, deepen existing ones or drill sidetracks off existing wells. As is described by Kellogg [2011] and Anderson et al. [2018], oil and gas companies usually do not own their own rigs but rent them from independent service providers together with specialized crews operating the rig. Since the number of total rigs and trained workers is approximately fixed in the short run – Anderson et al. [2018], in their model, treat rigs as a scarce resource that cannot be expanded – rental prices react strongly to changes in the oil or gas price [Anderson et al., 2018]. Usually, there are idle rigs available that are allocated on the spot market based on price. Kellogg [2014] notes, however that starting in 2004 rising natural gas and oil prices led to large fossil fuel companies locking up rigs in long term contracts. As a consequence, rig markets became extremely tight and therefore had a long unobservable wait list which lead to a disconnect between drilling decisions and drilling activity for some time. Therefore, Kellogg [2014] only uses drilling data through 2003 in his study, leaving later years out. Since the period in which I am interested – the shale gas expansion starting in the mid 2000s – lies after 2003, I cannot do the same. However, I want to state this fact as a caveat. One can see that until around 2008 the majority of active rigs was employed for natural gas drilling. After the financial crisis the number of active gas rigs somehow recovered before falling until 2016. During the same time, which was marked by high oil prices, the number of active rigs used for oil drilling increased steadily. The regime of high oil prices ended 2014 when OPEC+ increased production which started a price war on global oil markets. With some lag, oil drilling decreased dramatically and, together with gas drilling, reached its low in 2016 before



Figure 5: Active drilling rigs in the United States. Source: Baker Hughes Rig Count via Energy Information Administration (2022) 'Crude Oil and Natural Gas Drilling Activity' (monthly): https://www.eia.gov/dnav/ng/ng_enr_drill_s1_m.htm.

Notes: Lines mark takeoff of U.S. shale gas production (2007), oil price collapse due to increased OPEC+ production (2014) and opening of LNG liquefaction plant Sabine Pass I (2016).



Figure 6: International natural gas prices and price for crude oil (WTI). Source: Statistical Review of World Energy - BP via Our World in Data (2022) 'Natural gas prices' (annual): https://ourworldindata.org/grapher/natural-gas-prices.

Notes: Lines mark takeoff of U.S. shale gas production (2007), the Fukushima accident (2011) and the oil price collapse due to increased OPEC+ production (2014). Prices are given in current USD.

recovering somehow.

North American natural gas prices were not only low after 2008 in nominal terms or when compared to oil, but also in comparison to international natural gas prices. Figure 6 shows that after 2008 Canadian and U.S. natural gas prices dropped dramatically, while prices in Europe returned quickly to the levels experienced immediately before the crisis. LNG prices in South East Asia, however, rose even further and seemed to follow crude oil prices. Most likely, the reason for this differential between natural gas prices in Europe and Asia was increased LNG demand due to the Fukushima accident [Mu and Ye, 2018]. Even though natural gas prices in Europe and Asia fell after 2014, North American prices remained far below this level. But the shale gas boom enabled the United States not only to become a net exporter of natural gas and was accompanied by low prices but also facilitated a rapid increase in natural gas consumption.



Figure 7: Natural gas consumption in the United States. Source: Energy Information Administration (2022) 'Natural Gas Consumption by End Use' (annual): https: //www.eia.gov/dnav/ng/ng_cons_sum_dcu_nus_a.htm.

Figure 7 shows that until the shale boom natural gas consumption was rather stagnating with growing shares of electric power generation and shrinking shares of the industrial sector. After the shale boom, however, industrial gas consumption rose again while electric power generation turned to natural gas on a rapid scale. At the end of the sample, electric power plants were the single biggest consumer of natural gas in the United States. We will now turn towards a modelling framework that allows one to make some causal statements regarding the descriptive evidence presented in this section.

4 The Model

4.1 Econometric Set Up

I set up a structural vector autoregressive (SVAR) model of the form

$$A_0 y_t = c + \sum_{j=1}^p A_j y_{t-j} + A_x x_t + \epsilon_t,$$
 (1)

where y_t is a 4 × 1 vector consisting of active gas drilling rigs, natural gas marketed production, industrial production and the real price of gas in the United States. The $q \times 1$ vector x_t of exogenous variables consists of seasonally and trend adjusted heating and cooling degree days, dummies for extraordinary events or structural change and 11 monthly dummies to capture seasonal effects, c is an intercept, $A_0...A_p$ are 4 × 4 matrices, A_x is a 4 × q matrix and ϵ_t is a 4 × 1 vector consisting of Gaussian structural shocks. The individual structural shocks $\epsilon_{i,t}$ are assumed to be independent of each other and have unit variance, i.e.

$$(\epsilon_t) \stackrel{iid}{\sim} N(0, I_4)$$

The structural model (1) can also be written in reduced form as

$$y_t = d + \sum_{j=1}^p B_j y_{t-j} + B_x x_t + e_t, \quad (e_t) \stackrel{iid}{\sim} N(0, \Sigma),$$
(2)

where $B := [d, B_1, ..., B_p, B_x] = A_0^{-1}[c, A_1, ..., A_p, A_x]$ and $\Sigma = A_0^{-1}(A_0^{-1})'$. System (2) can then be expressed as a regression equation of the form

where $Y = [y_1, ..., y_T]$, D = 1 + 4p + q, column *i* of *Z* is $Z_i = [1, y'_{i-1}, ..., y'_{i-p}, x'_i]'$ and $E = [e_1, ..., e_T]$. Like Inoue and Kilian [2013] as well as Arias et al. [2018], I choose a Bayesian approach and propose a Gaussian-Inverse Wishart prior for the reduced form parameters. This implies the following prior distributions:

$$\Sigma \sim \mathcal{IW}_K(S_*, n) \tag{4}$$

and

$$\beta | \Sigma \sim N(\beta^*, V_\beta = V \otimes \Sigma), \tag{5}$$

where S_*, n, V as well as β^* (or equivalently B^* such that $\beta^* = vec(B^*)$) have to be pre-specified. As is described in [Kilian and Lütkepohl, 2018, 162-165], the Gaussian-Inverse Wishart prior is a natural conjugate prior, which means that the prior, the likelihood and the posterior are all from the same distributional family. In particular, one obtains the following posterior distribution:

$$\Sigma | y \sim \mathcal{IW}_K(S, \tau), \tag{6}$$

$$\beta | \Sigma, y \sim N(\bar{\beta}, \bar{\Sigma}_{\beta}), \tag{7}$$

where the parameters of the distribution are given as:

$$\begin{split} \bar{B} &= (B^* V^{-1} + YZ')(V^{-1} + ZZ')^{-1}, \\ \bar{\beta} &= vec(\bar{B}), \\ \bar{\Sigma}_{\beta} &= (V^{-1} + ZZ')^{-1} \otimes \Sigma, \\ S &= T\tilde{\Sigma} + S_* + \hat{B}ZZ'\hat{B} + B^*V^{-1}(B^*)' - \bar{B}(V^{-1} + ZZ')\bar{B}', \\ \tau &= T + n, \\ \hat{B} &= YZ'(ZZ')^{-1}, \\ \tilde{\Sigma} &= (Y - \hat{B}Z)(Y - \hat{B}Z)'/T. \end{split}$$

For the parameters of the prior distribution I stay agnostic and choose n = 0, $S_* = 0_{4\times 4}$, $V^{-1} = 0_{D\times D}$ and $B^* = 0_{4\times D}$. Thus the posterior means of the parameters boil down to

$$\mathbb{E}[B|Y] = \hat{B},$$

$$\mathbb{E}[\Sigma|Y] = (Y - \hat{B}Z)(Y - \hat{B}Z)'/(T - 4 - 1),$$

$$\mathbb{E}[\Sigma_{\beta}|Y] = (ZZ')^{-1} \otimes \mathbb{E}[\Sigma|y].$$

Therefore, the expected value of the posterior distribution of B is equal to the least squares estimate of B and the expected value of the posterior distribution of Σ is a consistent estimate for Σ in the frequentist sense under typical distributional assumptions. As a consequence, the Bayesian estimates of the parameters can also be interpreted from a frequentist viewpoint. It is inference in the sign identified structural model, for which the Bayesian interpretation of the estimates is needed.

4.2 Choice of Variables

I use four time series of endogenous variables (active rigs, gas production, industrial production and real price of gas) and two time series of exogenous variables (heating and cooling degree days) that are all available in monthly frequency from 1993M11 to 2020M01.³ As a measure of investment and drilling activity I include the number of active rigs currently used for gas drilling in the United States as provided by the Baker Hughes Rig Count.⁴ Since drilling is expensive and the economic viability of opening new wells is highly dependent on prices in the future [see Mu, 2019], active rigs serve as a barometer of expectations in the oil and gas industry. For the oil sector, Kellogg [2014] shows that price volatility influences drilling negatively and volatility measures derived from futures prices are a better predictor for drilling activity than backward-looking indicators that are based on past price variation. Therefore, the rig count measure should somehow serve to capture a speculative component in the gas market in the sense of Kilian and Murphy [2014] additional to technological progress in drilling and the exploitation of natural gas wells. But there is another major reason for including active rigs into the model: Recently, studies such as Anderson et al. [2018] have exploited engineering knowledge to explain the empirical fact that the short-term price elasticity of oil production is close to zero. According to them, production from existing wells is primarily determined by geological characteristics and as a consequence, producers mainly react to price changes via the drilling rate. For the natural gas industry, Mason and Roberts [2018] provide a similar result and report price elasticities of gas production from already existing wells that are small in magnitude.

The choice of the remaining endogenous series – gas production, economic activity and real price of gas – needs somewhat less explanation, because at least since Kilian [2009] these three variables lie at the center of many oil market SVAR models. As so often, however, the devil lies in the details. Following Arora and Lieskovsky [2014] and Wiggins and Etienne [2017], I use marketed U.S. natural gas withdrawals⁵ as a measure of gas production and real U.S. industrial production⁶ as a measure of economic activity. One could object to looking exclusively at the industrial sector, since industrial activity is only one part of economic activity and is continuously losing relevance compared to the government or service sector.⁷ While this might

³All series would be available until 2022M04. Since extraordinary events such as the Covid pandemic and the Texas energy crisis in 2021 provide additional challenges for modelling, I leave the last two years out of my analysis.

⁴Baker Hughes Rig Count via Energy Information Administration (2022) 'Crude Oil and Natural Gas Drilling Activity' (monthly): https://www.eia.gov/dnav/ng/ng_enr_drill_s1_m.htm.

⁵Energy Information Administration (2022): 'U.S. Natural Gas Marketed Production' (monthly): https://www.eia.gov/dnav/ng/hist/n9050us2M.htm

⁶FRED (2022) 'IPB50001N' (monthly): https://fred.stlouisfed.org/series/IPB50001N.

⁷Between 2009 and 2019 the industrial sector never contributed more than 20 percent to the GDP of the United States. Source: Statista (2022) 'Distribution of gross domestic product (GDP) across economic sec-

be true in general, Figure 7 reveals that for most of the sample the industrial sector was the single biggest consumer of natural gas in the United States, only getting surpassed by electric power generation in the last years. In contrast, gas consumption in the commercial, residential, and transportation sectors tends to be negligible. Additionally, the included weather variables should be responsible for a considerable share of the variation in residential and commercial natural gas usage. Moreover, in 2020 the industrial sector's electric energy consumption was about two thirds of residential electric energy consumption and roughly twenty percent lower than commercial electric power usage.⁸ Therefore, also a big chunk of all electric power generated is consumed by the industrial sector. This motivates using U.S. industrial production as a monthly measure of economic activity and not, for example, a factor containing information from multiple time series that are available on a monthly basis and relate to economic activity. One could ask why I do not use a measure of *qlobal* economic activity like the freight rate index proposed by Kilian [2009], as is common in oil market models. Like Wiggins and Etienne [2017] I argue that the North American natural gas market is essentially a local one. Li et al. [2014] show that between 1997 and 2011 there is evidence for convergence between European and Asian natural gas prices, while the North American natural gas price stands for itself. This should be especially true since the financial crisis of 2008, where North American natural gas prices unambiguously decoupled from international natural gas and crude prices (see Figure 6). For the real gas price the nominal Henry Hub price deflated by the U.S. consumer price index⁹ is used. The monthly Henry Hub price is provided by the Wall Street Journal and available for 1993M11-2014M02.¹⁰ From 1997M01 on a monthly Henry Hub series is made available by the EIA.¹¹ Like Arora and Lieskovsky [2014] I first use the Wall Street Journal series, but I switch to the EIA series when the former is not available anymore.

Finally, let's get to the exogenous variables. While Arora and Lieskovsky [2014] use residential natural gas demand as an endogenous variable in their VAR model, I do not include this measure since residential gas consumption is rather small and I want to keep the model parsimonious. Since variation of natural gas usage in the residential sector should come mainly from changes in temperature, I follow Brown and Yücel [2008] as well as Nick and Thoenes [2014] and include measures of temperature into the model. Heating or cooling degree days, as provided by the EIA¹², measure how cold or warm a certain location is. For an aggregate measure that can be used to assess national energy demand, local measures are population-weighted. Additionally to weather and seasonal dummies I also add dummy variables to exclude two extraordinary events and model one broken trend. The two events are the hurricanes Katrina in 2005M09 and Gustav in 2008M09. Katrina greatly reduced natural gas output, especially in the Gulf of Mexico, in an already tight market and led to a spike in natural gas prices [see DOE, 2006].

tors in the United States from 2000 to 2019' (annual): https://www.statista.com/statistics/270001/ distribution-of-gross-domestic-product-gdp-across-economic-sectors-in-the-us/.

⁸In 2020, the residential, commercial and industrial sector consumed 1,464.4, 1,287.4 and 959.1 Terawatt hours respectively. Source: Statista (2020) 'Electricity consumption in the United States in 2020, by sector': https://www.statista.com/statistics/560927/us-retail-electricity-consumption-by-sector/.

⁹FRED (2022) 'USACPIALLMINMEI' (monthly): https://fred.stlouisfed.org/series/ USACPIALLMINMEI.

¹⁰Wall Street Journal via FRED (2016) 'GASPRICE' (monthly): https://fred.stlouisfed.org/series/GASPRICE.

¹¹Energy Information Administration via FRED (2022) 'MHHNGSP' (monthly): https://fred.stlouisfed.org/series/MHHNGSP.

¹²Energy Information Administration (2022) 'Heating degree-days by census division' (monthly); 'Cooling degree-days by census division' (monthly): https://www.eia.gov/totalenergy/data/monthly/index.php

Three years later, Gustav also reduced natural gas production but was not followed by a price hike since the economic crisis of 2008 already reduced gas demand [see CRS, 2009]. The two hurricanes led to short-term reductions of gas production that are clearly extraordinary and violate the assumption of symmetric shocks. It is hard to imagine a productivity shock rising gas production multiple percentage points above the level expected in the previous month. I also use a broken trend dummy of the form [0, ..., 0, 1, 2, 3, ...] to account for a constant rise in gas production from the mid 2000s onwards. As is described by Kellogg [2011], Anderson et al. [2018], Wang and Krupnick [2015] and Mu [2019], among many others, from the early 2000s on innovations in hydraulic fracturing ('fracking') and horizontal drilling led to higher cost effectiveness and production levels in the oil and gas industry. It is not a straightforward choice to model this 'revolution' with exogenous broken trend dummies. First, since the increase in productivity in the oil and gas sector was not only the result of some breakthrough scientific innovations that can be clearly dated, but to a large extent the result of 'learning by drilling' [see Kellogg, 2011], one should interpret the shale boom as a smooth structural change rather than a clear break. Recently Demetrescu and Salish [2020] have shown that in the case of a smooth structural transition without a clear break date, ignoring the break and choosing the lag length of the VAR with the help of an information criterion is preferable to using dummy variables. One could also object to using exogenous dummies on economic grounds: High natural gas prices after 2000 helped making already existing fracking techniques economically viable and set an incentive for developing new techniques. In this sense, the shale gas boom must be partly seen as an endogenous outcome. However, due to the singularity of the event and the non-linearity as well as long delay times and horizons involved – as Wang and Krupnick [2015] highlight, efforts to exploit shale gas started multiple decades before 2000 – my SVAR model is clearly not able to disentangle the effects that led to the shale boom. The exogenous dummies have to be interpreted as an auxiliary construction to deal with this shortcoming. Regarding the econometric concerns, looking at the log-level of gas production plotted in Figure 8 makes it plausible to approximate part of the shale boom by a broken trend. The break in growth rates detected by the Bai and Perron [1998] procedure, 2005M05, is also historically plausible as start of the shale boom [Mu, 2019]. It has to be noted that the broken trend dummy does not alter the economic results qualitatively. Including it only makes the impulse response of gas supply to a gas supply shock less persistent and yields a more plausible historical decomposition of gas production.

Before modelling I apply unit root tests to the endogenous variables in log-levels. The ADF test with four lags never rejects the hypothesis of an unit root for all three specifications (no drift no trend, drift and no trend, drift and trend) on the 5 percent level. On the other hand, the KPSS test rejects the hypothesis of stationarity for natural gas production (specification 2 and 3) as well as for industrial production (specification 3) on the 1 percent level. For all other combinations the KPSS test does not reject the hypothesis on the 5 percent level. Since for the first differences the hypothesis of the ADF test is always rejected at the 1 percent level and the hypothesis of the KPSS test is never even rejected at the 10 percent level if we allow for a drift, I conclude that the endogenous variables are at most integrated of order one. Despite potential unit roots, I will estimate the model in log-levels. Since the Johansen eigenvalue test rejects the hypothesis of no cointegration at the 1 percent it is to conclude that there is at least one linear combination of active rigs, gas production, industrial production and the gas price in log-levels that is stationary. Estimating the model in differences would amount to throwing away information with regards to this equilibrium relationship. On the other hand, the potential

costs associated with estimating the model in log-levels are low. As [Kilian and Lütkepohl, 2018, 41-43] highlight, for integrated variables the parameters estimated by OLS are consistent and even standard Gaussian inference remains valid.¹³ Further, if one additional lag to the true lag order is added when estimating the model, the asymptotic variance covariance matrix of the slope parameters is of full rank, which justifies joint inference on the parameters. Of course there is some efficiency loss due to the additional parameters that have to be estimated. Also, in contrast to a stationary multivariate process corresponding to a stable VAR-system, an I(1) process does not have an $MA(\infty)$ representation. Therefore, the endogenous variables cannot be meaningfully approximated by a finite linear combination of past structural shocks as is done in the historical decomposition [see Kilian and Lütkepohl, 2018, 116-122]. There are, however, solutions to this problems to which I will come back later. Further, because a $MA(\infty)$ process is regular, its forecast error variance converges to its unconditional variance when the forecast horizon goes to infinity. This is not the case for integrated variables and therefore the forecast error variance decomposition does not converge. Thus, I will not deal with the variance decomposition of the forecast error for large horizons, in which I am not particularly interested anyway.

As a last step, the exogenous variables are transformed. Both heating and cooling degree days have a strong seasonal component and since seasonal dummies are already included in the VAR, the transformed variables should provide information on deviations from the typical seasonal temperature pattern. Also, even though the unit root tests do not detect any non-stationarity, knowledge of climate change makes it probable that the true data generating process is not stationary. This is also reflected by the descriptive fact that from 1980 to 2000 there are only three months with slightly more than 350 cooling degree days and from 2000 to 2020 there are eleven months with more than 350 cooling degree days and even two with more than 400. Since regressing an integrated variable on another integrated variable invalidates standard frequentist inference (spurious regression) it is appropriate to be cautious and therefore I choose a transformation that removes seasonality as well as a trend from the data and subtract the mean of the last five observations that were in the same month from each observation. Figure 8 shows all series as included in the VAR model.

4.3 Estimation of the Reduced Form

The reduced form of the model as given in (2) is estimated with OLS, choosing a lag length of 6 months despite the AIC suggesting only one lag. As a first reason, this is because Kilian and Zhou [2020] argue that traditional lag selection criteria often tend to underestimate the true lag order in finite samples and in particular for oil market models up to 24 monthly lags are needed to adequately capture the relationship between economic activity and oil market variables. Further, Kellogg [2014] reports that the engineering, permitting and rig contracting process leads to a gap of three months between the decision to drill a well and the start of the drilling process. Khalifa et al. [2017] also conclude that the impact of oil price changes on the number of active rigs has lags of multiple months. Even though both studies deal with the U.S. oil industry, these facts should also hold for the natural gas industry since drilling rigs and

¹³I am aware that the concerns regarding stationarity that are discussed here arise primarily from a frequentist approach and are not that relevant from a fully Bayesian perspective. However, since my reduced form estimates as well as Figure 10 are also interpretable from a frequentist viewpoint, I want to address these concerns.



Figure 8: Transformed variables that are included in the VAR model.

contractors are the same for both sectors.

Figure 9 plots the reduced form residuals which roughly look like the trajectories of a white noise process. The only exception is some evidence for volatility clustering when looking at the rig count and the gas price. I now apply diagnostic tests to the model residuals to check some assumptions previously imposed. Table 1 plots the result of the multivariate Breusch-Godfrey LM test for autocorrelation in the model residuals. For the standard autoregressive lag-length of 6 and 5 lags in the test equation, the hypothesis of no autocorrelation is rejected at the 1 percent level and for 14 lags in the test equation it is rejected at the 5 percent level. Increasing the number of autoregressive lags in the model reduces or even removes the evidence for residual autocorrelation. Using a lag number as high as 18, however, would greatly increase the dangers of overfitting and lead to a high number of parameters to be estimated. Therefore I will stick to the specification of 6 lags. As a compensation a robustness analysis is conducted, in which the model is estimated on the two sub-samples before and after the shale gas boom. This is motivated by the fact that the Breusch-Godfrey test does not reject the hypothesis of no autocorrelation if the model is fitted to the two individual sub-samples. As a last diagnostic exercise, the third and fourth moments of the model residuals are given in Table 2. Under the hypothesis of Gaussian errors the skewness should be 0 and the kurtosis should be 3. In the full sample, only the residuals of the number of active rigs most likely have nonzero skewness but



Figure 9: Reduced form residuals for the sample 1993M11-2020M01.

Sample:	1993M11-2020M01			1993M11-2005M04	2005M05-2020M01		
Lags Test	st $p = 6$ $p = 12$		p = 18	p = 6	p = 6		
5	0.001**	0.003**	0.438	0.604	0.812		
14	0.013^{*}	0.045^{*}	0.402	1	0.305		

 Table 1: Multivariate Breusch-Godfrey LM test for autocorrelated errors.

Notes: The table denotes p-values for different combinations of model lags (p) and lags used in the Breusch-Godfrey test. * p < 0.05, ** p < 0.01.

all variables seem to have some excess kurtosis. In the earlier sample there is no problem with skewness and only the residuals of rig and rpg have some excess kurtosis. In the later sample the rig count residuals, again, are problematic. Concluding, asides from the negative skewness of rpg and some excess kurtosis, the third and fourth moments of the model residuals do not look too problematic. In particular, the residuals of gas production seem to be symmetric after excluding hurricanes Katrina and Gustav using dummy variables.

4.4 Identification of the Structural Form

After estimating the reduced form, the structural form (1) is still to be identified since there exist infinitely many structural models corresponding to each reduced form. One can easily show that a structural model (1) corresponds to the reduced form (2) if and only if

	1993M11-2020M01			1993M11-2005M04			2005M05-2020M01					
	rig	gpd	ipd	rpg	rig	gpd	ipd	rpg	rig	gpd	ipd	rpg
Skew	-0.5^{**}	-0.1	-0.2	0.1	-0.1	0.3	0.3	-0.4	-0.7^{**}	-0.2	0.2	0.3
Kurt	4.8^{**}	3.6^{*}	3.7^{*}	4.0^{**}	4.0^{*}	2.9	2.8	4.0^{*}	5.2^{**}	3.5	3.5	2.8
J-B	52.8^{**}	5.7	7.8^{*}	13.2^{**}	5.9	2.3	2.2	8.2^{*}	48.6**	2.8	3.2	2.1

Table 2: Skewness and kurtosis of model residuals.

Notes: Skew denotes the standardized sample skewness, Kurt the standardized sample kurtosis and J-B the Jarque-Bera test statistic. Stars refer to the D'Agostino test (H0: no skewness), the Anscombe-Glynn test (H0: no excess kurtosis) and the Jarque-Bera test (H0: no skewness and no excess kurtosis). Abbreviations for variables: rig ... rig count, gpd ... gas production, ipd ... real industrial production, rpg ... real price of gas. * p < 0.05, ** p < 0.01.

$$A_0^{-1} = CQ,\tag{8}$$

$$[c, A_1, \dots, A_p, A_x] = A_0[d, B_1, \dots, B_p, B_x],$$
(9)

where $C = chol(\Sigma)$ and $Q \in \mathcal{O}(n) := \{Q \in \mathbb{R}^{n \times n} | QQ' = Q'Q = I_n\}$ is an arbitrary orthonormal matrix. To restrict the set of admissible structural models, one has to employ some identifying restrictions. As a first step, I require the impact effect matrix A_0^{-1} to match the sign/exclusion pattern given in Table 3. I assume that four orthogonal shocks – a drilling shock, a gas supply shock, an economic activity shock and a gas demand shock – span the space of the innovations to the four endogenous variables. A drilling shock leads to a rise in gas drilling. This is the only shock that can affect the number of active rigs contemporaneously, since for contractual, legal and technical reasons the decision to drill takes three months to result in actual drilling activities [see Kellogg, 2014]. With this assumption, the drilling shock as well as the first column of A_0^{-1} are point identified. When the elements of A_0^{-1} are denoted by a_{ij} one can compute $(\epsilon_{1,t})$ as the reduced form residual $(e_{1,t})$ divided by its standard deviation, a_{11} being equal to this standard deviation. The remaining elements $a_{1\bullet}$ can then be obtained by regressing $(\epsilon_{1,t})$ on the remaining model variables (or equivalently on the reduced form errors). Therefore there is no need to impose any sign restrictions on the first column of A_0^{-1} . The second structural shock, a gas supply shock, decreases gas production unexpectedly. As a consequence, the price of gas goes up and industrial production decreases. A gas supply shock can be an extreme weather event that decreases gas production for a short time or an unexpected change in well productivity. The third shock, an economic activity shock, raises industrial production on impact. This shock pushes gas prices up which triggers a rise in natural gas production. Lastly, a demand shock specific to the gas market rises gas prices and gas production while reducing industrial production. A gas demand shock can be thought of as a shift in preferences, representing changed heating or cooling behaviour or a higher amount of natural gas used for heating or power production instead of coal or oil due to environmental concerns. Another possible cause of gas demand shocks are substitution or spillover effects: Using data until 2007 Brown and Yücel [2008] report crude oil prices to be one of the main drivers of natural gas prices in the U.S. and for Germany Nick and Thoenes [2014] also conclude that oil prices are an important determinant of natural gas prices. Expanding the recursively identified model of Kilian [2009] to include a gas price variable, Jadidzadeh and Serletis [2017] estimate that

around 30 percent of the long run variance of the gas price is determined by fluctuations in the oil supply and oil specific demand. Therefore oil supply shocks, speculative oil demand shocks or rising oil prices due to a rise in foreign economic activity (which does not lead to higher U.S. industrial production) can find their way into the model as gas specific demand shocks. Since Nguyen and Okimoto [2019] report qualitatively similar responses of the gas market to oil demand and gas demand shocks, it is justifiable to not distinguish between the two in this paper's model. Of course the decoupling of oil and natural gas prices in the United States (see Figure 6) might have mitigated the importance of oil price shocks in recent years.

Not all structural models that satisfy the imposed impact sign restrictions are also economically plausible. Some may imply responses or elasticities which we know to be false. Kilian and Murphy [2012] therefore suggest using prior knowledge about such elasticities to further restrict the set of admissible solutions. Therefore I use microeconomic evidence provided in Mason and Roberts [2018], who study price elasticities of already existing natural gas wells in Wyoming. They report a price elasticity of gas production of 0.0262 for the whole sample. Splitting samples in four quartiles according to peak production rates gives no price elasticity that is different from zero for the first two quartiles. The third quartile has an elasticity of 0.0311 and the fourth quartile of 0.0646. To be prudent I use 0.065 as an upper bound for the price elasticity of oil supply and discard all model draws, for which a_{23}/a_{43} or a_{24}/a_{44} are above this threshold. Kilian and Murphy [2012] also restrict the impact response of economic activity to a gas demand shock to a small interval. The numerical value for this restriction is the hardest to motivate on empirical grounds so Kilian and Murphy choose an arbitrary value close to zero and conduct sensitivity tests in which they vary that threshold. I restrict the immediate impact of a gas demand shock on industrial production to lie in the interval (-0.004, 0), where 0.004 is half the standard deviation of the reduced form error of industrial production. As an additional guideline, one can use the theoretical result of Hamilton [2013] that the elasticity of output to a change in the real price of energy should equal the energy bill as a fraction of gdp times the price-elasticity of energy demand. This result only refers to direct, supply side effects and leaves aside changes in employment or capital formation which one can assume to operate with some lag. In 2008, the US natural gas bill was the highest in the sample with about 200 billion 2019 Dollars¹⁴, amounting to below 1.2 percent of gdp. Rubaszek et al. [2021] estimate a short term price elasticity of gas demand of -0.42. Even using the extremely cautious value of -1 for this elasticity would give a lower bound of roughly -0.0014 for the impact response of economic output to an unexpected price increase of 0.114one standard deviation of the reduced form error of the real gas price. This is about a third of the lower bound I use, since one has to account for the fact that industry is a disproportionally large natural gas user when compared to the total economy and natural gas price increases might be caused or accompanied by oil price increases. Further justification for a bound close to zero is given by Baumeister and Hamilton [2019] who use a prior mean of -0.05 for the oil price elasticity of economic activity and obtain a posterior of almost exactly zero. As a last identifying assumption, economic activity shocks are not allowed to lower industrial production in the first six months.

One could ask why I did not choose a recursive identification strategy like in Kilian [2009] for oil markets or Arora and Lieskovsky [2014] for natural gas markets. Indeed, as is pointed out above, there is strong evidence for the price elasticity of gas supply as well as the response of industrial production to a gas demand shock to be close to zero on impact. To assume

¹⁴Source: EIA, Today in Energy, September 9, 2021.

	Drilling	Gas Supply	Economic Activity	Gas Demand
	Shock	Shock	Shock	Shock
Rig Count	*	0	0	0
Gas Production	*	_	+	+
Ind. Production	*	_	+	_
Gas Price	*	+	+	+

Table 3: Impact sign and exclusion restrictions for A_0^{-1} .

Notes: Additional elasticity restrictions and dynamic sign restrictions are described in the text.

that these values are exactly zero would enable us to identify A_0^{-1} uniquely as the Cholesky decomposition of Σ . This would greatly simplify estimation, eliminate identification uncertainty and allow to use standard frequentist methods for inference. However, these values are, most likely, not exactly zero. Mason and Roberts [2018] show that the price elasticity of gas supply is close to zero but statistically significant different from zero for individual wells in Wyoming. Similarly, when looking at the European natural gas crisis after the Russian attack of Ukraine, one can see see that in principle industry can react quickly to rising gas prices. After a first surge of natural gas prices starting in March 2022, prices fell back to typical 2021 levels in May. After rising to March-levels again at the end of July, prices rose further to a new all time high at the end of August. Earlier in the year, Ruhnau et al. [2022] had already reported on- and off-switching of energy-intensive chemical plants in Germany in response to changes in natural gas prices. On August 23th Bloomberg reported that European plants had reduced one third of their ammonia output capacity [Gebre, 2022], on August 25th the same news provider reported a capacity reduction of 50 percent [de Sousa et al., 2022]. Of course, it is still possible that these capacity cuts were all decided in July or June, when natural gas prices started to rise again. However, since prices surpassed March-levels only one month before these extreme capacity reductions it is highly unlikely that all of these measures were planned months before. As an additional argument, chemical output can also to some extent be reduced without shutting down production capacity: Richardson [2022] reports that typical chemical plants can be operated at 60 percent of design capacity. Of course, the United States never experienced nearly as extreme energy price shocks in the observed sample period as Europe did in 2022. The events in Europe. however, at least point towards some short-term flexibility in the chemical industry and are sufficient to not believe in a zero-response of industrial production to energy price shocks in the same month. Therefore, a recursive identification strategy is not plausible. Nevertheless, I will later compare estimation results for the two identification strategies described here.

Now let us come to the actual identification procedure for the structural form. While the first shock is point identified, shocks 2-4 are only set identified. Since the probability of drawing Q such that CQ matches the zero restrictions in Table 3 is zero when one draws from an uniform distribution on $\mathcal{O}(4)$ the algorithm commonly used for SVAR models with just sign restrictions and no exclusion restrictions [see for example Kilian and Lütkepohl, 2018, 441] has to be adapted slightly. In particular, one can easily show that CQ matches the zero restrictions given above if and only if

$$Q = \begin{bmatrix} 1 & 0_{1\times3} \\ 0_{3\times1} & Q^*, \end{bmatrix}$$
(10)

where $Q^* \in \mathcal{O}(3)$. Therefore one can draw such a Q^* from an uniform distribution on $\mathcal{O}(3)$, build Q as in (10) and check if the sign restrictions are met. If this is the case, we can keep the draw. The detailed procedure is given in Algorithm 1. This algorithm exploits the fact that one can draw from $\mathcal{O}(n)$ by drawing each component of the $n \times n$ matrix J from a standard normal distribution and computing the QR-decomposition of J, [Q, R]. If one makes sure that each diagonal element of R is positive, the resulting Q is drawn from an uniform distribution on $\mathcal{O}(n)$.

Algorithm 1 Identify the set of admissible structural models **Input:** $B = [d, B_1, ..., B_p, B_x], C = chol(\Sigma), K$ (number of draws) **Output:** $k \leq K$ admissible structural models: $[c^j, A_0^j, A_1^j, ..., A_p^j, A_x^j], j \in \{1, ..., k\}$ 1: $j \leftarrow 1$ 2: for $i \in \{1, ..., K\}$ do draw a 3 × 3 matrix J with $J_v \sim N(0, I_3)$ for each column J_v of J 3: 4: let $[Q^* R]$ be the QR-decomposition of J for $v \in \{1, 2, 3\}$ do 5: if $R_{vv} < 0$ then 6: $Q_v^* \leftarrow -Q_v^*$ 7: end if 8: end for 9: build Q as in (10) 10:check if CQ satisfies impact sign restrictions, otherwise move to next draw 11: check if CQ satisfies elasticity restrictions, otherwise move to next draw 12:check if structural IRFs meet dynamic sign restrictions, otherwise move to next draw 13: $A_0^j \leftarrow (CQ)^{-1}$ 14: $[c^{j}, A_{1}^{j}, ..., A_{p}^{j}, A_{x}^{j}] \leftarrow A_{0}^{j}[d, B_{1}, ..., B_{p}, B_{x}]$ 15: $j \leftarrow j + 1$ 16:17: end for

Of course Algorithm 1 can be further optimized since it can get quite computationally burdensome if only a small fraction of all draws are admitted and therefore K has to be chosen to be quite large. A straightforward optimization is not calculating the full matrix CQ in the beginning but checking one element after the other by multiplying rows of C with columns of Q and discarding the draw as quickly as possible. Further, if the original Q^* is not admissible, one can switch signs of the columns and/or permutate the columns and check if the modified CQ now matches the requirement. This can be done because the resulting matrix Q^* will still be orthonormal. Also note that since at most one such permutation fulfils the sign restrictions the resulting matrices Q^* will still be uniformly drawn from $\mathcal{O}(3)$ if one stops trying different permutations once one combination is accepted.

4.5 Inference in the Structural Model

We now know how to identify the set of admissible structural models that correspond to a certain point estimate of reduced form parameters. Usually, however, one wants to use SVAR models to conduct inference, especially with regards to the impulse response functions. So how does one get point estimates of the impulse response function and corresponding highest density

or confidence bands? In many applications using sign restrictions, researchers report pointwise medians as well as pointwise quantile bands from a set of simulated impulse response functions as point estimate and highest density band, respectively. As Inoue and Kilian [2013] highlight, however, such pointwise statistics can be misleading. Pointwise median response functions do not need to match the behaviour of any simulated structural IRF. Likewise, pointwise error bands do not contain a certain fraction of admissible models. In general – not only in the context of sign restricted models – pointwise error bands for impulse response functions are too narrow to conduct joint inference. To address this problem there are several methods for joint inference on impulse response functions [see Lütkepohl et al., 2020, for an overview]. The typical solutions, however do not solve the problem which model to report as point estimate in the set identified case. An approach that provides broad enough confidence bands to contain a certain fraction of IRFs and one particular model as point estimate is provided by Inoue and Kilian [2013]. They suggest to report the modal model that maximizes the posterior density function of the IRF, $f(\Theta)$. After ordering all drawn solutions in decreasing order of their posterior density $f(\tilde{\Theta})$, one can compute the $(1 - \alpha)$ -quantile $f_{1-\alpha}$ of this value and define a $(1-\alpha)\%$ highest density region as $\mathcal{S}_{\alpha} = \{\Theta | f(\Theta) \geq f_{1-\alpha}\}$. Recently, Inoue and Kilian [2022] have generalized this approach using Bayesian loss functions, where the original procedure reporting the modal model is optimal using a Dirac-delta loss. Further, they show how to report the modal IRF in models identified by zero restrictions and exclusion restrictions using the algorithm of Arias et al. [2018]. I can use neither of the formulas without modification, since I draw Q not directly from $\mathcal{O}(4)$ but the smaller dimensional variable Q^* from $\mathcal{O}(3)$. In this case, it follows from Proposition 3 in the Appendix that when imposing an uniform prior on $Q^* \in \mathcal{O}(n-1)$, the posterior density of the structural IRFs can be computed as

$$f(\tilde{\Theta}) = 2^{(n^2 - n + 2)/2} |C|^{-np} \prod_{i=1}^{n} |c_{ii}|^{n-i+1} |I_{n-1} + \tilde{Q}^*|^{-(n-2)} \\ |\tilde{D}'_{n-1}(K' \otimes K'C')(I_{n^2} - L'_n L_n)(K \otimes CK)\tilde{D}_{n-1}|^{-1/2} f(\tilde{B}|\Sigma) f(\Sigma) f(s^*), \quad (11)$$

where n = 4, c_{ii} is the *i*th diagonal element of C and \tilde{Q}^* is just Q^* with the sign of the first column flipped if the determinant of Q^* is -1. The $(n-1)^2 \times (n-1)(n-2)/2$ matrix \tilde{D}_{n-1} satisfies $vec(S) = \tilde{D}_{n-1}veck(S)^{15}$ for a $(n-1) \times (n-1)$ skew-symmetric matrix S. Construction and central properties of this matrix are described in Magnus [1988]. Similarly L_n is $n(n+1) \times n^2$ such that $vec(L) = L'_n vech(L)$ for any lower triangular matrix L. Further, K is an $n \times (n-1)$ matrix defined as

$$K = \begin{bmatrix} 0_{1 \times (n-1)} \\ I_{n-1} \end{bmatrix}.$$

The last three terms in (11) are the posterior density of $\tilde{B} = [B_1, ..., B_p]$ conditional on Σ , the posterior density of Σ and the posterior density of $s^* = veck(S^*)$, where

$$S^* = I_{n-1} - 2(I_{n-1} + \tilde{Q}^*)^{-1}$$

 $^{^{15}}$ The operator *veck* stacks all below-diagonal elements of a matrix in a column vector, likewise *vec* stacks all columns of a matrix in a column vector and *vech* stacks all element on and below the diagonal in a column vector.

is a skew-symmetric matrix depending on Q^* . For $f(\tilde{B}|\Sigma)$ and $f(\Sigma)$ there are simple closedform solutions since they are densities of a multivariate normal distribution and a Gaussian-Inverse Wishart distribution, respectively. For the last factor, León et al. [2006] show that s^* has the density function

$$f(s^*) = \left(\prod_{i=2}^{n-1} \frac{\Gamma(i/2)}{\pi^{i/2}}\right) \frac{2^{(n-2)(n-3)/2}}{|I_{n-1} + S^*|^{n-2}}.$$
(12)

Therefore one can use Algorithm 2 to conduct Bayesian inference in the 4-dimensional SVAR model with sign and zero restrictions as set up above. For M draws from the reduced form parameters and K draws of orthogonal matrices Q^* for each reduced form draw, the algorithm returns $k \leq MK$ admissible structural models together with the posterior density of their IRFs. For inference I can now report the IRF corresponding to the modal model with the highest posterior density value as point estimate and plot the IRFs of $(1 - \alpha)100\%$ of all models as highest density region.

5 Results

5.1 Comparison of Identification Strategies

Before conducting inference, let us take a look at the IRFs of the structural models that correspond to the OLS estimate of the reduced form (2) using Algorithm 1. For comparison I also report the estimated IRFs corresponding to a recursive identification strategy. This amounts to assuming that $Q = I_4$ and therefore A_0^{-1} has the following form:

$$A_0^{-1} = \begin{bmatrix} * & 0 & 0 & 0 \\ * & * & 0 & 0 \\ * & * & * & 0 \\ * & * & * & * \end{bmatrix}.$$
 (13)

This is equivalent to assuming that a drilling shock immediately affects all variables, but is the only shock that affects the number of active rigs on impact – an assumption we already used and justified for the model above –, that the price elasticity of gas supply is zero and industrial production does not react to a gas demand shock in the same month. This basically means that we switch the identifying assumptions of Kilian and Murphy [2012] for the restrictions of Kilian [2009]. Note that some restrictions are similar in nature but the recursively identified model cannot be an admissible model corresponding to the sign restrictions [see Kilian and Lütkepohl, 2018, 421-424 for a discussion]. As I have argued above, the recursive identification assumptions are not plausible in my opinion. However, since many studies of the U.S. natural gas markets rely on such assumptions [Arora and Lieskovsky, 2014, Hou and Nguyen, 2018, Rubaszek and Uddin, 2020], it is interesting to compare the two strategies.

Figure 10 plots the resulting impulse response functions. Note that for the responses to the drilling shock, recursive identification and identification corresponding to Table 3 coincide. This is because the drilling shock is point identified by the assumption that it is the only shock affecting gas drilling on impact – an assumption that is in place in both models. We see that a drilling shock raises the number of active rigs in a hump shaped way. Also, gas production and

Algorithm 2 Conduct Bayesian inference in our SVAR model.

Input: $Y = [y_1, ..., y_T], Z = [Z_1, ..., Z_T], p, M, K$ **Output:** $k \leq MK$ admissible models with IRF posterior density: $[A_0^j, A_1^j, ..., A_p^j, f^j], j \in$ $\{1, ..., k\}$ 1: $j \leftarrow 1$ 2: compute reduced form posterior parameters as in (6) and (7)3: for $h \in \{1, ..., M\}$ do draw Σ from (6) 4: draw B conditional on Σ from (7) 5: $C \leftarrow chol(\Sigma)$ 6: $B \leftarrow [B_1, \dots, B_p]$ 7: for $i \in \{1, ..., K\}$ do 8: draw a 3 × 3 matrix J with $J_v \sim N(0, I_3)$ for each column J_v of J 9: let $[Q^* R]$ be the QR-decomposition of J 10: for $v \in \{1, 2, 3\}$ do 11: if $R_{vv} < 0$ then 12: $Q_v^* \leftarrow -Q_v^*$ 13:end if 14:end for 15:build Q as in (10) 16:check if CQ satisfies impact sign restrictions, otherwise move to next draw 17:check if CQ satisfies elasticity restrictions, otherwise move to next draw 18:check if structural IRFs meet dynamic sign restrictions, otherwise move to next draw 19: $\tilde{Q}^* \leftarrow Q^*$ 20: if $det(\hat{Q}^*) = -1$ then 21: $\tilde{Q}_1^* \leftarrow -\tilde{Q}_1^*$ 22: end if 23: $f^j \leftarrow f(\tilde{\Theta})$ as given in (11) 24: $A_0^j \leftarrow (CQ)^{-1}$ 25: $[\check{A_1^j},...,\check{A_p^j}] \leftarrow A_0^j[B_1,...,B_p]$ 26: $j \leftarrow j + \dot{1}$ 27:end for 28:29: end for

industrial production are raised slightly after 1-4 months. This delay is consistent with the fact that drilling and the fracturing process take some time and therefore natural gas production cannot be raised immediately after drilling has started. Let us now take a look at the responses to the remaining shocks. A gas supply shock decreases gas production sharply and has a persistent negative effect on natural gas output. This raises the real price of gas, increases gas drilling with a considerable lag and has no noticeable effect on industrial production. An economic activity shock increases industrial production persistently, increases the gas price and slightly increases gas production. One can see, however, that this supply response takes time and most likely depends on the increase in drilling that happens due to the rise in industrial production. Lastly, a gas demand shock rises the price of natural gas sharply but this response is considerably weakened after some months and entirely gone after one year. This shock decreases industrial production with the effect becoming stronger over time, which makes sense since adjustment to higher gas prices might take time. Gas production increases slightly and drilling also increases. It is noteworthy, however, that the response of drilling activity is less persistent than after an economic activity shock. We also see that the results corresponding to the recursive identification and to the sign restrictions do not differ too much qualitatively. The main difference is that the model with sign restrictions allows economic activity and gas supply shocks to have a stronger impact on gas prices and drilling than the recursively identified model. Also, in the sign restricted model gas demand shocks can have a stronger impact on industrial production. The magnitude of this effect, however, depends crucially on the limit to the impact response of industrial production to a gas demand shock. A tighter bound pushes the admissible responses in column 3 and 4 of row 3 and 4 further towards the recursively identified IRFs. We will come back to this discussion when talking about the forecast error variance decomposition.

5.2 Natural Gas Market Dynamics

As a next step one can compute the Bayesian point estimates and highest density bands for the impulse response functions that are computed using Algorithm 2. The results are plotted in Figure 11. The qualitative implications of the point estimates remain similar to above. Since they represent identification as well as parameter uncertainty and are jointly computed not pointwise, the 68% bands are quite broad. It is notable, however, that drilling, gas supply and economic activity shocks have a persistent effect on drilling activity, gas production and industrial production, respectively, that is significantly different from zero even after one year. Meanwhile, the effect of a gas demand shock on the gas price vanishes quickly and is only significantly different from zero for the first few months after the shock. Lastly, one can see that an economic activity shock raises gas prices and drilling activity significantly some months after the shock. As a robustness check, I also estimate the model on two sub-samples, 1993M11-2005M04 and 2005M05-2020M01, and compare the impulse response functions. The results can be seen in Figure 22 and 23 in the appendix. Most responses are qualitatively similar to the ones plotted in Figure 11. As a main difference, the response of gas production to a gas supply shock seems to be less persistent after splitting the sample. There are also some significant differences between the two samples: In the later time window, drilling shocks seem to have a stronger and more persistent effect on the number of active rigs and economic activity shocks, after some months, have a stronger effect on gas production than in the earlier sample. Furthermore, in



Figure 10: Admissible structural impulse response functions corresponding to OLS point estimate.

Notes: The black, dotted line corresponds to the recursive identification strategy. Colored lines are the set of admissible structural IRFs that satisfy the imposed sign restrictions.



Figure 11: Structural IRFs for the full sample 1993M11-2020M1. *Notes:* The black line is the modal model which maximizes the posterior density of the structural IRFs. Colored lines are the Bayesian 68% highest density region.

the first sample the effect of a drilling shock on gas production is significantly positive between month 2 and 4, while in the late sample the point estimate is negative and not significantly different from zero. Aside from that, the IRFs seem to match, which speaks for the robustness of the model results.

5.3 The Price Elasticity of Natural Gas Drilling

According to Figure 11, only natural gas price increases that are triggered by an economic activity shock lead to a significant response of the rig count. This is the case, even though the initial price change after a gas demand shock of one standard deviation is estimated to be bigger than after an economic activity shock of one standard deviation. This suggests, that the response of natural gas drilling to a change in the gas price depends on the shock that leads to the price change. Controlling for the size of the price change, these responses can be explicitly compared using price elasticities of drilling that can be computed as follows:

$$\kappa_h^k := \frac{\partial \log(rig_{t+h})}{\partial \log(p_t)} = \frac{\theta_{1k,h}}{\theta_{4k,1}}, \ k = 2, 3, 4, \tag{14}$$

where $\theta_{jk,h}$ is the element (j,k) in the structural IRF matrix of horizon h. It follows from this formula, that for each structural IRF three different price elasticities of drilling, κ_h^k , can be computed for every horizon h. After h months, a price change due to the structural shock k leads to a larger response of natural gas drilling than a price change due to shock l, if we have $\kappa_h^k - \kappa_h^l > 0$. Therefore, we can test differences in the magnitude of the drilling response by computing posterior distributions of the differences of the elasticities defined above, using the posterior distribution of the structural IRFs that we have already obtained. Figure 12 plots these distributions together with the fraction of all values in the posterior distribution that are negative. In the top panel, the differences of the price elasticities of drilling after an economic activity shock and a gas supply shock are plotted. We see that for a horizon of one quarter, more than 85 percent of the structural IRFs in the posterior distribution imply a stronger response of drilling after an economic activity shock than after a gas supply shock. However, this difference vanishes when we increase the horizon at which we look and after one year nothing suggests that the responses after an economic activity and after a gas supply shock are different. This picture is different when comparing economic activity shocks and gas demand shocks, which is done in the lower panel. After one quarter, more than 90% of all IRFs in the posterior distribution imply a stronger drilling response to price changes after economic activity shocks than after gas demand shocks. This fraction gets even bigger as we increase the horizon for which we compute elasticities. One year after the price change, the response of the rig count relative to the initial price change is bigger after an economic activity shock virtually for the whole posterior distribution. This makes sense, since the initial increase in prices after a gas demand shock has vanished by then, while after an economic activity shock prices are still increased significantly.

5.4 Forecast Error Variance Decomposition

To study the relative importance of different structural shocks for variation in the model variables, it is popular to calculate a forecast error variance decomposition. This is motivated by the fact that the h-step mean squared prediction error (MSPE) of an SVAR model is given by

$$\mathbb{E}[(y_{t+h} - \hat{y}_{t+h|t})(y_{t+h} - \hat{y}_{t+h|t})'] = \sum_{i=0}^{h-1} \Phi_i \Sigma \Phi'_i$$
$$= \sum_{i=0}^{h-1} \Theta_i \Theta'_i,$$

where Φ_i is the reduced form impulse response function of horizon i, Θ_i is the structural impulse response function of horizon i and $\hat{y}_{t+h|t}$ is the forecast of y_{t+h} conditional on information given at time t. The kth diagonal element of the resulting matrix is the h-step MSPE of variable



Figure 12: Posterior distributions of the differences of $\kappa_h = \partial \log(\operatorname{rig}_{t+h})/\partial \log(\operatorname{rpg}_t)$ by shock. Notes: This figure plots the posterior distribution of $\kappa_h^i - \kappa_h^j = \theta_{1i,h}/\theta_{4i,1} - \theta_{1j,h}/\theta_{4j,1}$ for two structural shocks *i* and *j*.

k that can be written as

$$MSPE^{k}(h) = \sum_{i=0}^{h-1} (\theta_{k1,i}^{2} + \dots + \theta_{kn,i}^{2})$$
$$= \sum_{j=1}^{n} (\theta_{kj,1}^{2} + \dots + \theta_{kj,h-1}^{2}), \qquad (15)$$

where $\theta_{kj,i}$ denotes element (k, j) of Θ_i . Equation 15 shows that $MSPE^k(h)$ can be decomposed in contributions from different shocks, using the structural IRFs. In particular, one can make use of the fact that

$$1 = \frac{\theta_{k1,1}^2 + \dots + \theta_{k1,h-1}^2}{MSPE^k(h)} + \dots + \frac{\theta_{kn,1}^2 + \dots + \theta_{kn,h-1}^2}{MSPE^k(h)}.$$
 (16)

More information on the forecast error variance decomposition can be found in [Kilian and Lütkepohl, 2018, 113-116]. Figure 13 plots the decomposition given in (16) for each n = 1, ..., 4and horizon h = 0, ..., 14. For consistency, the model corresponding to the modal IRF displayed in Figure 11 is used. One can see that short-term variation in the number of active rigs is mainly driven by drilling shocks – which is not a big surprise since drilling takes some time to respond to unexpected shocks other than drilling shocks. After some months, however, economic activity shocks gain in importance and also gas supply as well as gas demand shocks explain a small fraction of the forecast error. Variation in gas production is mostly determined by gas supply shocks with gas demand shocks being the second most important factor. For industrial production, economic activity shocks are the most important driver of variation in the forecast error. What is perhaps surprising is the relative importance of gas demand shocks. Since natural gas is an important fuel for industry, natural gas price fluctuations that happen because of changed preferences or extraordinary weather events in regions with disproportionally high natural gas usage for heating and electricity generation, might affect industrial production. Also, oil price shocks working their way into the model as gas demand shocks via substitution effects may do the same. However, I must note that the share of the forecast error of industrial production that is due to gas demand shocks is mainly governed by the bound for the immediate impact of a gas demand shock on industrial production. There are good reasons to impose such a bound, but it is hard to argue for a particular number. Therefore I have set the relatively prudent lower bound of half a standard deviation of the reduced form residuals of industrial production times minus one. A narrower band reduces the relative importance of gas demand shocks for industrial production. Figure 26 and 27 in the Appendix plot impulse response functions for alternative bounds to the immediate response of industrial production to a gas demand shock. In particular, I restrict this response to lie in the interval (-b, 0), where b is now not half but a third of the standard deviation of the reduced form residual of industrial production. This reduces not only the immediate impact of a gas demand shock on industrial production but also – and to a lesser extent – of an economic activity shock to the gas price. The qualitative implications of the impulse response functions, however, remain unaltered. Concerning the forecast error variance decomposition, gas demand now explains a larger (but still minor) fraction of the forecast error of gas production and a larger share of the forecast error of the gas price, while the influence on industrial production is reduced dramatically. At the same time, gas supply shocks now explain a larger fraction of the forecast error of the gas price – which is still at most 25 percent – while the share of the



Figure 13: Forecast error variance decomposition for the full sample 1993M11-2020M1. *Notes:* Forecast error variance decomposition corresponds to the impulse response function with maximal posterior density.

economic activity shock shrinks. Another robustness check is to look at the results for the two sub-samples previously used for the impulse response analysis. These are presented in Figure 24 and 25 in the Appendix. In the earlier sample, drilling shocks seem to explain a smaller share of the forecast error corresponding to the number of active rigs than in the late sample. Economic activity shocks, in the late sample, seem to be responsible for a large part of variation in gas production after some months. When it comes to the effects of gas demand shocks on industrial productions, there is a similar issue as in the full sample – we know, however, that this result crucially depends on the bound for the effect of the gas demand shock on industrial production.

Some results concerning the forecast error variance decompositions are sensitive to the choice of b, for which I have no precise empirical estimate. Therefore I have to stick to the results that are robust to the choice of b. In particular the model shows that drilling and gas supply shocks explain the largest fraction of the forecast error of active rigs and gas production, respectively. After some months, however, economic activity shocks explain a considerable fraction of the forecast error of the rig count and gas demand shocks also play a considerable role for gas production. The forecast error of industrial production is only to a small extent explained by drilling or gas supply shocks. Lastly, gas supply shocks can only explain a minor fraction of the forecast error of the gas price. Instead, the gas price is mainly driven from the demand side, but we are not in a position to exactly estimate the relative importance of gas demand and economic activity shocks.

Historical Decomposition in the Presence of Unit Roots

Impulse response analysis and forecast error variance decomposition provide some insights into the inner workings of U.S. natural gas markets. However, to assess the relative importance of certain shocks in different historical periods they are not an appropriate tool. Instead, this can be done using the historical decomposition of endogenous variables. Given an AR(p) system $B(L)(y_t) = (e_t)$, where $(e_t) \sim WN(\Sigma)$ and $B(L) = I_n - B_1 - \dots - B_p$ is a stable $n \times n$ lag-polynomial, the process (y_t) has an $MA(\infty)$ representation

$$y_t = \sum_{j \ge 0} \Phi_j e_{t-j} = \sum_{j \ge 0} \Theta_j \epsilon_{t-j}, \tag{17}$$

where Φ_j is the reduced form IRF of horizon j, Θ_j is the structural IRF $\Phi_j A_0^{-1}$ and (ϵ_t) is the series of structural shocks such that $\epsilon_t = A_0 e_t$. Since (y_t) is an $MA(\infty)$ process it is regular and therefore the effects of past shocks vanish over time. This motivates the approximation

$$y_{t+h} \approx \sum_{j=0}^{h} \Theta_j \epsilon_{t+h-j},$$
(18)

that converges in probability to the true value of y_{t+h} if h goes to infinity. Therefore, one can decompose the kth endogenous variable the following way:

$$y_{k,t+h} \approx \sum_{j=0}^{h} [\theta_{k1,j} \epsilon_{1,t+h-j} + \dots + \theta_{kn,j} \epsilon_{n,t+h-j}]$$

= $\sum_{i=1}^{n} \sum_{j=0}^{h} \theta_{ki,j} \epsilon_{i,t+h-j} =: \sum_{i=1}^{n} \hat{y}_{k,t+h}^{i}.$ (19)

Intuitively, $\hat{y}_{k,t+h}^{i}$ represents the contribution of shock *i* to the endogenous variable *k* at time t + h. If this value is positive, shock i has a positive influence on variable k and the other way around. Another way to think of $y_{k,t+h}^i$ is as a counterfactual value for $y_{k,t+h}$ for the case of all past shocks but i having been zero and shock i having been as observed. Even though the insights that such a historical decomposition can provide are potentially useful to study natural gas markets, I cannot simply use formula (19) to compute it. The reason is that since my system is not stable, the effects of structural shocks cannot be expected to vanish over time and therefore approximation (18) is not valid. This makes it necessary to adapt the formula and its interpretation slightly. Given endogenous variables $[y_{-p}, ..., y_0, ..., y_T]$, exogenous variables $[x_0, ..., x_T]$, an estimate for the reduced form parameters B as defined in (2) and a structural parameter matrix \hat{A}_0 like in (1) one can compute estimates for the structural shocks $\mathcal{E} = [\hat{\epsilon}_0, ..., \hat{\epsilon}_T]$ as

$$\mathcal{E} = \hat{A}_0 (Y - \hat{B}Z). \tag{20}$$

Afterwards one defines $\hat{\epsilon}_t^i = [0, ..., 0, \hat{\epsilon}_{i,t}, 0, ..., 0]'$ for $i \in \{1, ..., n\}$ and recursively computes the decomposition

$$\hat{y}_t^i = \hat{d} + \sum_{j=1}^p \hat{B}_j \hat{y}_{t-j}^i + \hat{B}_x x_t + \hat{A}_0^{-1} \hat{\epsilon}_t^i,$$
(21)

with starting values $[\hat{y}_{-p}^i, ..., \hat{y}_0^i] := [y_{-p}, ..., y_{-1}]$. Note that this is not a decomposition in the sense provided above for the case of stable AR(p) systems without exogenous variables. Neither does \hat{y}_t^i only contain past structural errors *i*, nor is the sum of all decompositions an approximation of the true value y_t . The first point is due to the fact that the starting values contain all kind of structural shocks whose effects cannot be expected to vanish over time. Instead, \hat{y}_i^* has to be interpreted as the counterfactual value of y_t for the case of all structural shocks other than shock i having been zero from time zero on. The second point raises the question how to present the historical decomposition such that the relative contributions of different structural shocks can be interpreted qualitatively. For $MA(\infty)$ processes, a natural comparison is the constant zero-function, since in the absence of any shocks the process converges to this value. In the case of my model, zero is no meaningful comparison value. Also, comparing y_t^i to the true value y_t does not tell us if shock i does exert a positive or negative influence on the endogenous variables at time t. Suppose, for example, that at time t all structural shocks have a positive influence on $y_{k,t}$. Then $\hat{y}_{k,t}^i$ lies below $y_{k,t}$ for every *i*. To get a series of meaningful comparison values one has to set $\hat{\epsilon}_t^0 = [0, ..., 0]$ and compute \hat{y}_t^0 using the recursive procedure given in (21). This is the counterfactual value of y_t for the case of no shocks from time zero on. If shock i has a positive influence on $y_{k,t}$ then $\hat{y}_{k,t}^i$ lies above $\hat{y}_{k,t}^0$. Further, this comparison value has the nice property that the differences $(\hat{y}_{k,t}^i - \hat{y}_{k,t}^0)$ sum up to $(y_{k,t} - \hat{y}_{k,t}^0)$, which somehow justifies the term decomposition.

5.6 Decomposition of Drilling Activity

The historical decompositions of gas production and industrial production are not particularly interesting for my questions but consistent with common intuition and major historical events. For the remaining two variables, however, the method described above provides valuable insights. The historical decomposition of the number of active rigs is presented in Figure 14. The top panel plots the true number of active rigs (grey) and the comparison value $y_{1,t}^0$ (black), while in the lower four panels the differences of the real rig count to the comparison value as well as of the counterfactual values to the comparison values, i.e. $(y_{1,t} - y_{1,t}^0)$ and $(y_{1,t}^i - y_{1,t}^0)$, are displayed. Despite continuously higher gas production, active drilling rigs seem to experience a downward trend since 2005, which could be due to more effective drilling and higher well productivity. One can see that in the years before 2008, economic activity was more important in driving up drilling activity to all-time heights than were drilling shocks. After the great recession, however, when economic activity, gas demand and gas prices were low, drilling remained extraordinarily stable. Egan [2016] notes this fact and explains it with an increase in drilling and well productivity, which matches the estimated importance of drilling shocks in this period. Afterwards, in 2016, oil and gas drilling in the United States reached a 70 year low. This is most likely due to low business confidence in an environment of extremely cheap natural gas and oil, bankruptcies and vast excess capacities. At the end of 2016, finally, drilling rates increase rapidly. As Blackmon [2016] explains, this was not only because of higher natural gas prices but also due to huge efficiency gains in drilling leading to lower costs per drilled well, new discoveries of shale gas reservoirs in the Permin basin and extraordinary cheap drilling equipment. Also, as Blackmon [2016] notes, in 2016 more than 200 oil and gas companies went bankrupt, which cleared out excess capacities and most likely is responsible for the positive



Figure 14: Historical decomposition of the number of active rigs.

Notes: The comparison value $y_{1,t}^0$ is plotted in black, the grey line in the top panel is the true rig count $y_{1,t}$ as included in the model. The remaining panels present the historical decomposition series that were computed using (21) minus the comparison value, i.e. $(y_{1,t}^i - y_{1,t}^0)$ (color), as well as the true rig count minus the comparison value, i.e. $(y_{1,t}^i - y_{1,t}^0)$ (color), as well as the true rig count minus the comparison value, i.e. $(y_{1,t}^i - y_{1,t}^0)$ (color), as well as the true rig count minus the comparison value, i.e. $(y_{1,t}^i - y_{1,t}^0)$ (grey). This is done to highlight the fact that $\sum_i (y_{1,t}^i - y_{1,t}^0) = (y_{1,t} - y_{1,t}^0)$ justifies the term decomposition.

contribution of supply shocks to drilling activity.

5.7 Disentangling Historical Natural Gas Price Fluctuations

Now let us get to the main point of this exercise, the historical decomposition of the natural gas price as presented in Figure 15. One can see that the comparison value increases until 2005 and then decreases. This changed trend can be attributed to the shale gas boom. However, there is much that cannot be explained by the deterministic parts of the time series. Before 2000, natural gas prices were unusually low. This was mainly due to gas specific demand shocks. A possible explanation for this would be extremely cheap oil after the Asian financial crisis – in November 1998, the nominal price for one Barrel of WTI was the lowest since 1986.¹⁶ Some years later, in 2000 and 2001, natural gas prices spiked. The model correctly attributes the biggest share of this increase to rising gas demand. The reason for the surge in natural gas prices was the California energy crisis, where artificial tightening of electric energy supply and market manipulation led to a surge in demand for natural gas. As Wilson [2002] points out, at this time

¹⁶Source: FRED (2022) 'DCOILWTICO' (monthly): https://fred.stlouisfed.org/series/DCOILWTICO.



Figure 15: Historical decomposition of the gas price.

Notes: The comparison value $y_{4,t}^0$ is plotted in black, the grey line in the top panel is the true rig count $y_{4,t}$ as included in the model. The remaining panels present the historical decomposition series that were computed using (21) minus the comparison value, i.e. $(y_{4,t}^i - y_{4,t}^0)$ (color), as well as the true rig count minus the comparison value, i.e. $(y_{4,t}^i - y_{4,t}^0)$ (color), as well as the true rig count minus the comparison value, i.e. $(y_{4,t}^i - y_{4,t}^0)$ (color), as well as the true rig count minus the comparison value, i.e. $(y_{4,t}^i - y_{4,t}^0)$ (grey). This is done to highlight the fact that $\sum_i (y_{4,t}^i - y_{4,t}^0) = (y_{4,t} - y_{4,t}^0)$ justifies the term 'decomposition'.

around 50 percent of Californian electricity generation capacity was driven by natural gas with expensive natural gas power plants acting as the swing producer. Therefore, when Californian electricity demand rose, natural gas demand rose disproportionally. Additionally, the high electricity prices led to older, less efficient natural gas power generators to be re-activated. As a result, from May 2000 to August 2001 natural gas demand for electricity generation in the system of the Pacific Gas and Electric Company was 41 percent higher than in the twelve months before. When the effects of this crisis vanished, the fall of gas prices was reinforced by a drop in industrial demand due to the economic crisis in the aftermath of the Dotcom bubble. Additionally, falling oil prices may explain some portion of the fall in gas demand. In 2005, natural gas prices reached their record level in the sample. The reasons are supply disruptions due to the hurricanes Katrina and Rita. I have already accounted for the former with a dummy. However, a small portion of the remaining rise in prices is correctly estimated to be due to a gas supply shock. Additionally, around this time economic activity shocks begin to exert a positive influence on prices which they continued to do until the financial crisis of 2008. Shortly before the great recession, natural gas prices spiked again. According to CRS [2009], a rise in demand from the power generation (+2.8%) and residential (+3%) sector in an already tight market has to be made responsible for this, while industrial demand rose only 0.3% in



Figure 16: Contributions of trend and different shocks to natural gas price collapse after 2008. *Notes:* Solid lines mark changes in log real natural gas prices in two time windows. Colored bars represent contributions of fitted trend and different structural shocks to this price change that were computed using (22).

2008. Additionally, the price of WTI hit a record high of almost 140 USD/Barrel in June 2008. Again, the model identifies demand shocks as the main driver for this spike in prices. Wiggins and Etienne [2017] report similar results. According to them, the exact origin of the price spike is somehow unsettled. However, they review some potential explanations that are all in line with my definition of a gas specific demand shock. After this peak, the gas price was in free fall from 2008 to 2012. Note that it are demand side factors and not supply changes due to the shale boom that are mainly responsible for this price drop. This is not self-evident, since economists and market analysts have attributed this price collapse largely to the shale boom [Wiggins and Etienne, 2017]. Hausman and Kellogg [2015] for example estimate that demand side factors are responsible for less than 20 percent of the fall in natural gas prices between 2007 and 2013, while the shale gas revolution is argued to be responsible for the remainder. My model results challenge the predominant role of supply side factors. To assess the relative contributions of the trend and the structural shocks to the change of some variable one can compute:

$$y_{k,t+h} - y_{k,t} = (y_{k,t+h} - y_{k,t+h}^{0}) + (y_{k,t+h}^{0} - y_{k,t}^{0}) + (y_{k,t}^{0} - y_{k,t})$$
$$= (y_{k,t+h}^{0} - y_{k,t}^{0}) + \sum_{i=1}^{n} [(y_{k,t+h}^{i} - y_{k,t+h}^{0}) + (y_{k,t}^{0} - y_{k,t}^{i})], \qquad (22)$$

where $(y_{k,t+h}^0 - y_{k,t}^0)$ is the part of the change that can be explained by the trend (as well as seasonal and weather effects) and the remaining summands are the contributions of different structural shocks *i*. The first two bars of Figure 16 plot the contributions of different factors to the gas price collapse after its peak in 2008M06. Ten months after this date the natural gas price reached a local minimum. According to the historical decomposition implied by the modal IRFs, more than two thirds of this collapse can be attributed to demand side factors. After some recovery, the natural gas price fell even deeper in 2012M04. It is perhaps more surprising than for the first time span that still demand side effects accounted for most of the drop in prices. Figure 17 shows the posterior distribution of the fraction of the natural gas price drop that was due to demand side shocks (economic activity shocks and gas demand shocks). For the period 2008M06-2009M09 the whole 68% (equal-tailed) confidence interval is above 75 percent and for 2008M06-2012M04 it is still above 50 percent. Thus we can conclude that for the natural gas price collapse after 2008, demand side factors were more important than the shale gas revolution. The third bar of Figure 16 plots the contributions of different shocks and the trend to the change in natural gas prices between June 2007 and 2013, which is roughly the same time window that Hausman and Kellogg [2015] look at. While it is obvious that the trend component pulled prices down strongly, supply shocks actually dragged it up somehow. Thus, demand side shocks – of which economic activity shocks are the most important – are responsible for more than 45 percent of the drop in natural gas prices, which is more than double the fraction that Hausman and Kellogg [2015] estimate. However, Figure 17 shows that for this time interval the estimate of the contribution of demand side shocks to the drop in natural gas prices is fairly imprecise. Therefore, the value of Hausman and Kellogg [2015] – which is estimated using an instrument depending on variation in weather conditions and therefore also has a high variance – lies within the 68% confidence interval of my estimate. Additionally, there are some differences between my approach and the one of Hausman and Kellogg [2015]: They compare yearly average prices of 2007 and 2013, whereas I look at monthly prices. Furthermore, they estimate long-run price elasticities of natural gas supply and demand to construct (long run) supply and demand curves, while I use the historical decomposition implied by my SVAR model. This makes the two estimation results not entirely comparable, but my estimates suggest that demand-side shocks may have played an important role for the observed decline in natural gas prices also in the period 2007-2013.

The origins of the economic activity shocks in 2008 are straightforward. The gas demand shock is most likely the result of low consumer confidence and a severe reduction of household wealth. Further, the oil price – itself in free fall – which used to pull the gas price upwards in the years before, now became decoupled from the latter as was pointed out before. Also, this was the time when the United States almost completely withdrew from global LNG markets before returning as a net supplier in 2016. After the financial crisis, economic activity and gas demand held the natural gas price below its trend until about 2014 before playing a positive role again. One possible explanation for the persistent rise in gas demand is the growing importance of natural gas for electric power generation, as can be seen in Figure 7. In 2016, natural gas surpassed coal as the most important source of electric power in the U.S. [Walton, 2017]. In the same year, however, the natural gas price fell to its lowest level in the sample. As Egan [2016] reports, in this year the U.S. energy industry was struggling with excess supply. After 2016 prices rose again, which, according to Walton [2017] was due to a combination of supply disruptions and stronger demand. Additionally, the bankruptcy of more than 200 oil and gas companies [see Blackmon, 2016] may have also cleared out some excess supply, which would be represented by a supply shock as it is the case in our decomposition. All in all, the historical decomposition of natural gas prices seems to closely match the prevailing explanation of many events. One possible difference, however, is the natural price collapse after 2008, which my model largely attributes to demand side factors instead of the shale gas boom.



Figure 17: Posterior distribution of the fraction of the natural gas price drop that can be attributed to demand side shocks (economic activity and gas demand shocks) in three time intervals.

Notes: Red line represents fraction that corresponds to the modal IRF, dotted lines are endpoints of the 68% equal-tailed confidence intervals. Plotted are distributions of the random variable $\frac{100 \cdot [(y_{4,t+h}^3 - y_{4,t+h}^0) + (y_{4,t+h}^0 - y_{4,t+h}^0) + (y_{4,t}^0 - y_{4,t+}^0)]}{y_{4,t+h} - y_{4,t}} | Y, Z, \text{ for three combinations of } t \text{ and } h.$

6 Conclusion

To study natural gas market dynamics in the United States I have set up a structural vector autoregressive model that includes the number of active drilling rigs, natural gas production, industrial production and the real price of gas. To identify the structural form I employ a set of both, exclusion and sign restrictions. For conducting inference I adapt the procedure of Inoue and Kilian [2013, 2018] that was originally designed for purely sign identified models. Many of my results are similar to the ones of existing studies on the U.S. gas markets conducted, among others, by Arora and Lieskovsky [2014], Wiggins and Etienne [2017] and Rubaszek et al. [2021]. In particular, I show that gas demand and economic activity shocks account for the biggest part of variation of the natural gas price and supply-side shocks are comparatively negligible.

As a new contribution, I present evidence for the fact that the response of drilling to unexpected changes in the price of gas depends on the source leading to the price change. Drilling reacts sluggishly to the other model variables and price changes due to gas specific demand shocks are typically large but vanish quickly. Therefore, economic activity shocks that change prices less but more persistently raise drilling significantly and persistently, while gas demand shocks do not. In particular, the price elasticity of drilling after an economic activity shock is significantly larger than after a gas demand shock. Since previous studies like Anderson et al. [2018] and Mason and Roberts [2018] have found out that the rate of drilling is the most important determinant of oil and natural gas supply, this result is highly relevant for policymakers that want to expand domestic energy production. In the case of the United States, a considerable amount of the record high levels of gas drilling before 2008 can be attributed to economic activity shocks. After 2008, technological advances in drilling prevented the number of active rigs to collapse despite exceptionally low natural gas prices.

Expanding natural gas production is usually not an end in itself. Although national security of supply and geopolitical autonomy have recently become the focus of renewed attention, governments typically hope that greater supply will lead to cheaper energy for national businesses and the electorate. It is therefore worthwhile to disentangle the reasons for the U.S. natural gas price collapse after 2008. Hausman and Kellogg [2015] attribute this fall to the shale gas revolution – a view they seem to share with most economists and analysts [Wiggins and Eti-

enne, 2017] – and estimate that demand side factors are responsible for less than 20 percent of the total price drop between 2007 and 2013. However, my model highlights the importance of demand side factors. I estimate – albeit with a wide confidence interval – that between June 2007 and June 2013 more than 45 per cent of the price decline was due to demand-side shocks. When looking at a narrow time window from the June 2008 price high until September 2009. according to my model more than two thirds of the price drop can be attributed to economic activity and gas demand shocks. Even when comparing the natural gas price of June 2008 to the approximately 85 percent lower value of April 2012, I conclude that the contribution of demand side factors to this drop was greater than that of supply side factors. The two latter results even apply to the entire 68% equal-tailed confidence interval. The importance of demand side shocks for the gas price drop after 2008 has been previously highlighted by Wiggins and Etienne [2017] and Rubaszek et al. [2021] who also make use of an SVAR model. However, they do not compute and present the contributions of different shocks to the change in the price level over different periods. Thus, this studies do not challenge the estimate of Hausman and Kellogg [2015] explicitly, even though calculations to do so would have been straightforward in principle. This evidence of the importance of demand side shocks after 2008 should dampen the hopes of governments working on their own shale gas booms. As follows from Wang and Krupnick [2015] and Mu [2019], most countries lack one or many of the peculiarities that led to the roaring success of shale gas development in the United States.¹⁷ However, even with a comparatively rapid expansion of shale gas production, no one should expect a price decline similar to that seen in the United States after 2008.

¹⁷Examples for this are water availability, deep capital markets, a dynamic service-sector industry and the specific market structure of the oil and gas industry in the United States.

7 List of Figures

1	U.S. (shale) gas production
2	U.S. natural gas trade
3	U.S. LNG trade and capacity
4	Oil and natural gas prices
5	U.S. active drilling rigs
6	International natural gas prices and oil price
7	U.S. natural gas consumption
8	Model variables
9	Reduced form residuals
10	IRFs: comparison of recursive identification and sign restrictions
11	IRFs and error bands
12	Comparison of price elasticities of drilling by shock
13	Forecast error variance decomposition
14	Historical decomposition: rig count
15	Historical decomposition: gas price
16	Contributions to gas price collapse after 2008
17	Contributions of demand side shocks to natural gas price drop
18	Rotations vs. Sub-Rotations
19	Mapping Reduced Form to IRFs
20	IRFs: Original Inoue and Kilian [2013] Oil Model
21	IRFs: Replication of the Inoue and Kilian [2013] Oil Model
22	IRFs: early sample 1993M11-2005M04 58
23	IRFs: late sample 2005M04-2020M01
24	FEVD: early sample 1993M11-2005M04
25	FEVD: late sample 2005M04-2020M01
26	IRFs: robustness check
27	FEVD: robustness check

8 Bibliography

- S. T. Anderson, R. Kellogg, and S. W. Salant. Hotelling under pressure. Journal of Political Economy, 126(3):984–1026, June 2018.
- J. E. Arias, J. F. Rubio-Ramírez, and D. F. Waggoner. Inference based on structual vector autoregressions identified with sign and zero restrictions: Theory and applications. *Econometrica*, 86(2):685–720, March 2018.
- V. Arora and J. Lieskovsky. Natural gas and u.s. economic activity. *The Energy Journal*, 35 (3):167–182, July 2014.
- J. Bai and P. Perron. Estimating and testing linear models with multiple structural changes. *Econometrica*, 66(1):47–78, January 1998.
- C. Baumeister and J. D. Hamilton. Sign restrictions, structural vector autoregressions, and useful prior information. *Econometrica*, 83(5):1963–1999, September 2015.
- C. Baumeister and J. D. Hamilton. Structural interpretation of vector autoregressions with incomplete identification. *American Economic Review*, 109(5):1873–1910, May 2019.
- D. Blackmon. The oil and gas situation 2016 year in review, Dec. 2016. URL https://www.forbes.com/sites/davidblackmon/2016/12/20/ the-oil-and-gas-situation-2016-year-in-review/?sh=4e6a15187c65.
- S. P. A. Brown and M. K. Yücel. What drives natural gas prices? *The Energy Journal*, 29(2): 45–60, 2008.
- CRS. Natural gas markets: An overview of 2008. Congressional Research Service Report for Congress, 2009.
- A. de Sousa, S. Gebre, and S. Treloar. Europe's widening fertilizer crisis threatens food supplies, August 2022. URL https://www.bloomberg.com/news/articles/2022-08-25/ yara-to-further-cut-european-ammonia-production-due-to-gas-spike.
- M. Demetrescu and N. Salish. (structural) var models with ignored changes in mean and volatility. *Working Paper*, November 2020.
- DOE. Impact of the 2005 hurricanes on the natural gas industry in the gulf of mexico region. United States of America Department of Energy Report, 2006.
- M. Egan. U.s. oil and gas rigs fall to 70-year lows, Mar. 2016. URL https://money.cnn.com/ 2016/03/11/investing/oil-gas-rig-count-baker-hughes-70-year-low/.
- S. Gebre. Europe's fertilizer industry shrinks again as energy chaos bites, Aug. 2022. URL https://www.bloomberg.com/news/articles/2022-08-23/ europe-s-fertilizer-industry-shrinks-again-as-energy-chaos-bites#xj4y7vzkg.
- J. D. Hamilton. What is an oil shock? Journal of Econometrics, 113(2):363–398, April 2003.
- J. D. Hamilton. *Handbook on Energy and Climate Change*, chapter Oil Prices, Exhaustible Resources, and Economic Growth, pages 29–63. Edward Elgar Publishing, 2013.

- C. Hausman and R. Kellogg. Welfare and distributional implications of shale gas. Brookings Papers on Economic Activity, pages 71–125, 2015.
- C. Hou and B. H. Nguyen. Understanding the us natural gas market: A markov switching var approach. *Energy Economics*, 75:42–53, 2018.
- IGU. Wholesale gas price survey: A global review of price formation mechanisms. Technical report, International Gas Union, 2021.
- A. Inoue and L. Kilian. Inference on impulse response functions in structural var models. Journal of Econometrics, 177(1):1–13, 2013.
- A. Inoue and L. Kilian. Corrigendum to "inference on impulse response functions in structural var models". *Journal of Econometrics*, 209(1):139–143, March 2018.
- A. Inoue and L. Kilian. Joint bayesian inference about impulse responses in var models. *Journal* of *Econometrics*, 231(2):457–476, December 2022.
- A. Jadidzadeh and A. Serletis. How does the u.s. natural gas market react to demand and supply shocks in the crude oil market? *Energy Economics*, 63:66–74, 2017.
- R. Kellogg. Learning by drilling: Interfirm learning and relationship persistence in the texas oilpatch. Quarterly Journal of Economics, 126(4):1961–2004, November 2011.
- R. Kellogg. The effect of uncertainty on investment: Evidence from texas oil drilling. American Economic Review, 104(6):1698–1734, June 2014.
- A. Khalifa, M. Caporin, and S. Hammoudeh. The relationship between oil prices and rig counts: The importance of lags. *Energy Economics*, 63:213–226, 2017.
- L. Kilian. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–1069, June 2009.
- L. Kilian and H. Lütkepohl. *Structural vector autoregressive analysis*. Themes in modern econometrics. Cambridge University Press, 2018.
- L. Kilian and D. P. Murphy. Why agnostic sign restrictions are not enough: understanding the dynamics of oil market var models. *Journal of the European Economic Association*, 10(5): 1166–1188, October 2012.
- L. Kilian and D. P. Murphy. The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, 29(3):454–478, April/May 2014.
- L. Kilian and X. Zhou. The econometrics of oil market var models. *Federal Reserve Bank of Dallas Working Paper*, March 2020.
- C. A. León, J.-C. Massé, and L.-P. Rivest. A statistical model for random rotations. *Journal of Multivatiate Analysis*, 97(2):412–430, 2006.
- R. Li, R. Joyeux, and R. D. Ripple. International natural gas market integration. *The Energy Journal*, 35(4):159–179, October 2014.

- H. Lütkepohl, A. Staszewska-Bystrova, and P. Winker. Constricting joint confidence bands for impulse response functions of var models - a review. *Econometrics and Statistics*, 13:69–83, 2020.
- J. R. Magnus. *Linear structures*. Oxford University Press, 1988.
- C. F. Mason and G. Roberts. Price elasticity of supply and productivity: An analysis of natural gas wells in wyoming. *The Energy Journal*, 39(1):79–100, 2018.
- X. Mu. The Economics of Oil and Gas. Agenda Publishing, 2019.
- X. Mu and H. Ye. Towards an integrated spot lng market. *The Energy Journal*, 39(1):211–234, January 2018.
- B. H. Nguyen and T. Okimoto. Asymmetric reactions of the us natural gas market and economic activity. *Energy Economics*, 80:86–99, 2019.
- S. Nick and S. Thoenes. What drives natural gas prices? a structural var approach. *Energy Economics*, 45(C):517–527, 2014.
- J. Richardson. Europe's gas crisis: the implications for global chemicals, July 2022. URL https://www.icis.com/asian-chemical-connections/2022/07/ europes-gas-crisis-the-implications-for-global-chemicals/.
- M. Rubaszek and G. S. Uddin. The role of underground storage in the dynamics of the us natural gas market: A threshold model analysis. *Energy Economics*, 87, 2020.
- M. Rubaszek, K. Szafranek, and G. S. Uddin. The dynamics and elasticities on the u.s. natural gas market. a bayesian structural var analysis. *Energy Economics*, 103, 2021.
- O. Ruhnau, C. Stiewe, J. Muessel, and L. Hirth. Gas demand in times of crisis. the response of german households and industry to the 2021/22 energy crisis. *Working Paper*, 2022.
- S. Shakya, B. Li, and X. Etienne. Shale revolution, oil and gas prices, and drilling activities in the united states. *Energy Economics*, 108, 2022.
- R. Walton. Eia: Natural gas prices in 2016 were the lowest in 2017. URL https://www.utilitydive.com/news/ nearly 20years, Jan. eia-natural-gas-prices-in-2016-were-the-lowest-in-nearly-20-years/434091/.
- Z. Wang and A. Krupnick. A retrospective review of shale gas development in the united states: What led to the boom? *Economics of Energy and Environmental Policy*, 4(1):5–17, 2015.
- S. Wiggins and X. L. Etienne. Turbulent times: Uncovering the origins of us natural gas price fluctuations since deregulation. *Energy Economics*, 64:197–205, 2017.
- J. F. Wilson. High natural gas prices in california, 2000-2001: Causes and lessons. Journal of Industry, Competition and Trade, 2(1-2):39–57, June 2002.

A Mathematical Appendix

Here I will prove some statements made in Sections 4.4 and 4.5. The first Proposition is a central and well known result for sign identified SVARs.

Proposition 1. Suppose the unknown matrix $A \in \mathbb{R}^{n \times n}$ is of full rank and AA' = S. Then a matrix B is a possible solution for A if and only if B = CQ, where C = chol(S) and $Q \in O(n)$.

Proof. Suppose we have such a B. Then we have

$$BB' = CQQ'C' = CC' = S.$$
(23)

Now let B satisfy BB' = S. Then we have

$$B = C \overbrace{C^{-1}B}^{Q:=}, \tag{24}$$

where $\tilde{Q} \in \mathcal{O}(n)$ because

$$\tilde{Q}\tilde{Q}' = C^{-1}BB'(C^{-1})' = C^{-1}S(C^{-1})' = C^{-1}CC'(C')^{-1} = I_n,$$

and similarly $\tilde{Q}'\tilde{Q} = I_n$.

The second result proves the claim that in this paper's model the matrix Q must have the form (10) in a more general case.

Proposition 2. Suppose the unknown matrix $A \in \mathbb{R}^{n \times n}$ is of full rank and is block lower triangular, i.e. there exist $n_1 + \ldots + n_k = n$ such that A consists of k^2 block matrices $A_{ij} \in \mathbb{R}^{n_i \times n_j}$ with $A_{ij} = 0_{n_i \times n_j}$ for j > i. Further, A satisfies AA' = S. Then a matrix B is a possible solution for A if and only if B = CQ, where C = chol(S) and $Q = diag(Q_{n_1}, \ldots, Q_{n_k})$ is block-diagonal with $Q_i \in \mathcal{O}(n_i)$.

Proof. Following Proposition 1 it is clear that B must have the form

$$B = CQ, \tag{25}$$

with $Q \in \mathcal{O}(n)$. Since C is lower triangular, C^{-1} is lower triangular and in particular block lower triangular as defined above. Therefore $Q = C^{-1}B$ is the product of two block lower triangular matrices. Using the same partition as for A, one can compute each of the k^2 blocks of Q as

$$Q_{ij} = (C^{-1}B)_{ij} = \sum_{s=1}^{k} (C^{-1})_{is} B_{sj}.$$
(26)

Now suppose j > i. Then for s = 1, ..., j - 1 the matrix B_{sj} consists only of zeros and for s = j, ..., k the same is true for matrix $(C^{-1})_{is}$. Therefore in this case the right side of (26) is a sum of zero matrices and Q_{ij} consists of zeros only. Thus Q is also block lower triangular. As

a next step we will prove that using the same partition, every orthonormal matrix Q that is block lower triangular is in fact block diagonal with orthonormal matrices as diagonal entries. For k = 1 this statement is trivial. Now suppose it holds for a certain k and take an arbitrary matrix $Q \in \mathcal{O}(n)$ with a partition corresponding to $n_1 + \ldots + n_{k+1} = n$ such that Q is block lower triangular. We can then partition this matrix into

$$Q = \begin{bmatrix} Q_{11} & 0_{n_1 \times n_{\bullet}} \\ Q_{\bullet 1} & Q_{\bullet \bullet}, \end{bmatrix},$$
(27)

where Q is of dimension $n_1 \times n_1$, $Q_{\bullet 1}$ is of dimension $n_{\bullet} \times n_1$ with $n_{\bullet} = n_2 + ... + n_{k+1}$ and $Q_{\bullet \bullet}$ is of dimension $n_{\bullet} \times n_{\bullet}$. Note that $Q_{\bullet \bullet}$ is block lower triangular with k^2 blocks. We can then use the fact that $Q \in \mathcal{O}(n)$:

$$QQ' = \begin{bmatrix} Q_{11} & 0_{n_1 \times n_{\bullet}} \\ Q_{\bullet 1} & Q_{\bullet \bullet} \end{bmatrix} \begin{bmatrix} Q'_{11} & Q'_{\bullet 1} \\ 0_{n_{\bullet} \times n_{1}} & Q'_{\bullet \bullet} \end{bmatrix} = \begin{bmatrix} Q_{11}Q'_{11} & Q_{11}Q'_{\bullet 1} \\ Q_{\bullet 1}Q'_{11} & Q_{\bullet 1}Q'_{\bullet 1} + Q_{\bullet \bullet}Q'_{\bullet \bullet} \end{bmatrix} \stackrel{!}{=} \begin{bmatrix} I_{n_{1}} & 0_{n_{1} \times n_{\bullet}} \\ 0_{n_{\bullet} \times n_{1}} & I_{n_{\bullet}} \end{bmatrix}.$$
(28)

One can derive similar conditions from $Q'Q \stackrel{!}{=} I_n$. It follows immediately that $Q_{11} \in \mathcal{O}(n_1)$. Thus left-multiplying $Q_{11}Q'_{\bullet 1}$ by Q'_{11} yields $Q_{\bullet 1} = 0_{n_{\bullet} \times n_1}$ from which we obtain $Q_{\bullet \bullet} \in \mathcal{O}(n_{\bullet})$. Therefore $Q_{\bullet \bullet}$ is an orthonormal, block lower triangular matrix with k^2 blocks which, by induction assumption, is block diagonal with orthonormal matrices as diagonal entries. It follows that the statement holds also for k + 1 and therefore for all $k \geq 1$, which finishes the proof.

Setting $n_1 = 1$, $n_2 = 3$ and k = 2 it follows from Proposition 2 that for the model of this paper it must hold that $A_0^{-1} = CQ$ with $C = chol(\Sigma)$ and Q being as given in (10). The general statement, however, has also a nice geometric interpretation. Suppose we have an *n*dimensional vector of reduced form errors (e_t) with variance-covariance matrix Σ and (ϵ_t) is an *n*-dimensional vector of uncorrelated structural shocks with unit variance. Let us assume that both variables are connected via

$$A(e_t) = (\epsilon_t),\tag{29}$$

where A^{-1} is an $n \times n$ matrix that is block lower triangular, as is given in Proposition 2. Then it follows from the same statement that $A^{-1} = CQ$ with $Q = diag(Q_1, ..., Q_k), Q_i \in \mathcal{O}(n_i)$. Thus we get

$$(u_t) := C^{-1}(e_t) = Q(\epsilon_t).$$
 (30)

Geometrically, (u_t) is the vector of uncorrelated errors obtained by using the Gram Schmidt orthonormalization procedure. This vector is equal to the identified 'structural' shocks if a recursive identification strategy is used. Since Q is block-diagonal, we get the following identity for every $i \leq k$:

$$\begin{bmatrix} (u_{m_i+1,t}) \\ \vdots \\ (u_{m_i+n_i,t}) \end{bmatrix} = Q_i \begin{bmatrix} (\epsilon_{m_i+1,t}) \\ \vdots \\ (\epsilon_{m_i+n_i,t}) \end{bmatrix},$$
(31)

where $m_i := n_1 + ... + n_{i-1}$. Since Q_i is a regular square matrix, the orthogonal errors $(u_{m_i+1,t}), \dots, (u_{m_i+n_i,t})$ span the subspace of the structural errors $(\epsilon_{m_i+1,t}), \dots, (\epsilon_{m_i+n_i,t})$. In other words, imposing block lower triangularity on A^{-1} allows one to identify subspaces in the space of all structural shocks – which, in SVAR analysis, is assumed to equal the space of the reduced form shocks – that are spanned of $n_i \leq n$ of the structural shocks. It is only within this subspaces that rotations need to be applied to transform the errors (u_{it}) to the true structural shocks (ϵ_{it}) . The reduced dimensionality of the rotations is matched by the reduced dimensionality of the matrices Q_i , n_i , when compared to n. Suppose for example that we are in the three dimensional case with k = 2, $n_1 = 1$ and $n_2 = 2$. Then we will let the direction of $(e_{1,t})$ untouched, project $(e_{2,t})$ on $[(e_{1,t})]^{\perp}$, $(e_{3,t})$ on $[(e_{1,t}), (e_{2,t})]^{\perp}$ and rotate the resulting vectors $(u_{2,t})$ and $(u_{3,t})$ in the 2-dimensional plane defined by $[(e_{1,t})]^{\perp}$ to obtain all three structural shocks up to scale. Figure 18 visualizes this procedure and its difference to 'ordinary' rotations. Finding three uncorrelated structural shocks that span the space of the observed reduced form shocks amounts to finding an orthogonal coordinate system of said space (the rows of A_0^{-1} are then the coordinates of the reduced form shocks with respect to this coordinate system). In general, this task has no unique solution. The Figure plots the coordinate system that is obtained by the Cholesky decomposition, i.e. (e_1) is spanned by (ϵ_1) , (e_2) is a combination of the first two coordinate vectors and the third coordinate vector is orthogonal to the first two. If we then rotate the whole coordinate system, we end up with another set of orthogonal vectors of unit length, i.e. another coordinate system. If we impose the restriction that the first column of A_0^{-1} is of the form [* 0 0], however, the possible rotations are severely restricted. If the direction of (ϵ_1) would be different from the direction of (e_1) the imposed zero restrictions were violated. Therefore we may only rotate the first axis, which is equivalent to rotating (ϵ_2) and (ϵ_3) in $[(\epsilon_2), (\epsilon_3)] = [(\epsilon_1)]^{\perp} = [(e_1)]^{\perp}$. Proposition 2 also covers two corner cases: If k = 1we are left with the same identification problem as in the case of classical sign restrictions. In the case of k = n we have $n_i = 1$ for every $i \leq k$ and every such subspace is one dimensional. Thus the structural shocks are point identified up to scale. If some of the blocks have dimension $n_i = 1$, these shocks are point identified, while the remaining shocks are only 'block identified' - as I call it - and additional sign, elasticity or shape restrictions may be used to narrow down the set of potential solutions.

There remains the question how to conduct inference in the block identified case. In the case of classical sign restrictions, i.e. k = 1, Inoue and Kilian [2013, 2018] prove that in a Bayesian framework one can explicitly derive the posterior density of the structural impulse response functions $\tilde{\Theta}$ as

$$f(\tilde{\Theta}) = \left| \frac{\partial vec(\tilde{\Theta})}{\partial [vec(B)' \ vech(C)' \ s']} \right|^{-1} \left| \frac{\partial vech(\Sigma)}{\partial vech(C)} \right| f(B|\Sigma) f(\Sigma) f(s), \tag{32}$$

where $S = I_n - 2(I_n + \tilde{Q})^{-1}$, s = veck(S) and the operators vec, vech and veck stack the columns, lower triangular and sub-diagonal elements of a matrix, respectively. \tilde{Q} is Qmanipulated to satisfy $|\tilde{Q}| = 1$. For the remainder of the section, when I write Q I will refer to \tilde{Q} . The matrix B is defined as $[B_1, ..., B_p]$. León et al. [2006] show that for an uniform distribution on $\{Q \in \mathcal{O}(n) ||Q| = 1\}$ the variable s has the density function



Notes: This figure displays random variables interpreted as vectors in a Hilbert space (the square root of the covariance being the inner product). Viewed through this lens, the length of a vector is its standard deviation, the cosine of two vectors is their correlation coefficient and orthogonality in the geometric and probabilistic sense are equivalent. The plotted coordinated system is the set of structural shocks obtained using the Cholesky decomposition (or equivalently the Gram-Schmidt procedure). Possible rotations of the coordinate system that are allowed under the identification assumptions are drawn on the axes.

$$f(s) = \left(\prod_{i=2}^{n} \frac{\Gamma(i/2)}{\pi^{i/2}}\right) \frac{2^{(n-1)(n-2)/2}}{|I_n + S|^{n-1}}.$$
(33)

With the posterior density being explicitly computable one can choose the modal model that has the highest posterior density as point estimate. Also, by plotting impulse responses for the $100(1 - \alpha)\%$ of all drawn models with the highest posterior density, one can construct joint confidence bands for the impulse responses that avoid the problems of pointwise bands. The result of Inoue and Kilian [2013], however, depends on the change-of-variable method which holds because there is a one-to-one mapping between [vec(B)', vech(C)', s'] and the first p + 1 structural IRFs, $[\theta_0, \theta_1, ..., \theta_p] = [A_0^{-1}, \Phi_1 A_0^{-1}, ..., \Phi_p A_0^{-1}]$, where Φ_i are the reduced form impulse responses at horizon *i*. In the case of sub-rotations, however, one has to adapt the formula above. This is because we draw not from $\mathcal{O}(n)$ but from smaller dimensional subspaces $\mathcal{O}(n_i)$. Therefore a mapping between [vec(B), vech(C), s] and $\tilde{\Theta}$ cannot be bijective, since additional linear restrictions on Q lead to the dimension of veck(S) (and of Q) being smaller than n(n-1)/2. I will derive an adaption of the formula of Inoue and Kilian [2013, 2018] for the special case of k = 2, $n_1 = 1$ and $n_2 = n - 1$ such as in the model set up in this paper. In this case, it follows from the dimensional arguments made above that there is a one to one mapping

$$\Phi: [vec(B)' \ vech(C)' \ (s^*)'] \mapsto vec(\tilde{\Theta}), \tag{34}$$



Figure 19: One-to-one Mapping from Reduced Form and Sub-Rotation to first p IRFs. *Notes:* Dots represent unrestricted elements that can be chosen freely from \mathbb{R} and matrix elements without dots are functions of the unrestricted elements. One can verify by counting that the two objects in the grey boxes have the same dimension.

where $s^* = veck(S^*)$ and $S^* = I_{n-1} - 2(I_{n-1} + Q^*)^{-1}$. The matrix Q^* is the one used to construct Q like in (10). The matrix S^* is skew-symmetric and thus all the information about the rotation represented by Q^* is contained in the lower-diagonal elements of Q^* . Figure 19 illustrates that indeed Ψ is a mapping between two vector spaces of the same dimension. It follows that one can compute the posterior density of the structural IRFs as

$$f(\tilde{\Theta}) = \overbrace{\left|\frac{\partial vec(\tilde{\Theta})}{\partial [vec(B)' \ vech(C)' \ (s^*)']}\right|^{-1}}^{I} \overbrace{\left|\frac{\partial vech(\Sigma)}{\partial vech(C)}\right|}^{II} \overbrace{f(B|\Sigma)}^{III} \overbrace{f(\Sigma)}^{V} \overbrace{f(s^*)}^{V}.$$
(35)

Terms II-IV are exactly like in Inoue and Kilian [2018], term V is given as (33), where one just has to substitute S^* for S and n-1 for n. Only deriving I needs some tedious computations. Before doing that, however, I should introduce some notation: Let D_n be the $n^2 \times n(n+1)/2$ duplication matrix such that $vec(M) = D_n vech(M)$ for any $n \times n$ symmetric matrix M. Further, \tilde{D}_{n-1} is the $(n-1)^2 \times (n-1)(n-2)/2$ matrix designed to satisfy $vec(S^*) = \tilde{D}_{n-1}s^*$. Similarly L_n is $n(n+1)/2 \times n^2$ such that $vec(M) = L'_n vech(M)$ for any lower triangular matrix M. The $n \times 1$ vector e_i^n consists of zeros with only the *i*th element being equal to one. Lastly, let us define the $n \times (n-1)$ matrix K as

$$K = \begin{bmatrix} 0_{1 \times (n-1)} \\ I_{n-1} \end{bmatrix}.$$
(36)

Now the following proposition that justifies using equation (32) can be proven:

Proposition 3. Suppose we have an $n \times n$ SVAR-model, where only the first of n structural shock is allowed to have an immediate impact on the first variable and we draw from an uniform distribution on $\mathcal{O}(n-1)$ to identify the model via sub-rotations as outlined above. Then the joint posterior density of the structural impulse response functions can be computed as

$$f(\tilde{\Theta}) = 2^{(n^2 - n + 2)/2} |C|^{-np} \prod_{i=1}^{n} c_{ii}^{n-i+1} |I_{n-1} + Q^*|^{-(n-2)} |\tilde{D}'_{n-1}(K' \otimes K'C')(I_{n^2} - L'_n L_n)(K \otimes CK)\tilde{D}_{n-1}|^{-1/2} f(B|\Sigma) f(\Sigma) f(s^*).$$
(37)

Proof. As first and major step, we will take the derivative of $vec(\Theta)$ with respect to $vec(\Phi)'$, vech(C)' and s'. Note that $Q^* = 2(I_{n-1} - S^*)^{-1} - I_{n-1}$ and therefore

$$dQ^* = 2(I_{n-1} - S^*)^{-1} dS (I_{n-1} - S^*)^{-1}.$$
(38)

It follows that

$$vec(dQ^*) = 2(((I_{n-1} - S^*)^{-1})' \otimes (I_{n-1} - S^*)^{-1})vec(dS)$$

= 2((((I_{n-1} - S^*)^{-1})' \otimes (I_{n-1} - S^*)^{-1})\tilde{D}_{n-1}ds. (39)

Now comes the step where we depart from Inoue and Kilian [2018]. We can express Q as

$$Q = e_1^n (e_1^n)' + KQ^* K'.$$
(40)

It follows that

$$d(CQ) = d(C)Q + Cd(KQ^*K')$$

= $d(C)Q + 2CK(I_{n-1} - S^*)^{-1}dS^*(I_{n-1} - S^*)^{-1}K'$ (41)

$$d(\Phi CQ) = d(\Phi)CQ + \Phi d(C)Q + 2\Phi CK(I_{n-1} - S^*)^{-1}dS^*(I_{n-1} - S^*)^{-1}K', \qquad (42)$$

and further

$$d(vec(CQ)) = (Q' \otimes I_n)L'_n d(vech(C)) + 2(I_n \otimes C)vec(K(I_{n-1} - S^*)^{-1}dS^*(I_{n-1} - S^*)^{-1}K')$$

= $(Q' \otimes I_n)L'_n d(vech(C)) + 2(I_n \otimes C)J_Q ds^*$ (43)

$$d vec(\Phi CQ) = (Q'C' \otimes I_{np})d(vec(\Phi)) + (Q' \otimes \Phi)L'_nd(vech(C)) + (I_n \otimes \Phi C)J_Qds^*,$$
(44)

where $J_Q := (K((I_{n-1} - S^*)^{-1})' \otimes K(I_{n-1} - S^*)^{-1})\tilde{D}_{n-1}$. Note that the final results for the differentials are just as in Inoue and Kilian [2018], with the only difference being the definitions of J_Q and s^* . While in Inoue and Kilian [2018] J_Q is a square matrix in our case it is not, since we have to match the different dimensions of S^* and C. This difference in dimensions will remain the major challenge in the following proof, since, in general, square matrices are easier to handle than non-square matrices.

Remembering that $\tilde{\Theta} = [(CQ)' \ (\Phi CQ)']'$ it follows immediately from (43) and (44) that

$$J_1 := \frac{\partial vec(\Theta)}{\partial [vech(C)' \ (s^*)' \ vec(\Phi)]} = \begin{bmatrix} (Q' \otimes I_n)L'_n & (I_n \otimes C)J_Q & 0_{n^2 \times n^2 p} \\ (Q' \otimes \Phi)L'_n & (I_n \otimes \Phi C)J_Q & (Q'C' \otimes I_{np}) \end{bmatrix}.$$
(45)

Because this matrix is block lower triangular its determinant is equal to the product of the determinants of the matrices on the diagonal:

$$|J_1| = |(Q' \otimes I_n)L'_n (I_n \otimes C)J_Q||(Q'C' \otimes I_{np})|$$

= |(Q' \otimes I_n)L'_n (I_n \otimes C)J_Q||C|^{np}. (46)

Now we can re-write

$$[(Q' \otimes I_n)L'_n \ (I_n \otimes C)J_Q] = (Q' \otimes I_n)Z, \tag{47}$$

with

$$Z = [L'_n \ 2(F \otimes G)\tilde{D}_{n-1}], \tag{48}$$

$$F = QK(I_{n-1} - (S^*)')^{-1} = \frac{1}{2}K(I_{n-1} + Q^*) = K(I_{n-1} - S^*)^{-1},$$
(49)

$$G = CK(I_{n-1} - S^*)^{-1} = CF.$$
(50)

As a next step we can make use of the equality

$$|(Q' \otimes I_n)Z| = |Z| = \sqrt{|Z|^2} = \sqrt{|Z'Z|}.$$
 (51)

Because of

$$Z'Z = \begin{bmatrix} L_n \\ 2\tilde{D}'_{n-1}(F'\otimes G') \end{bmatrix} \begin{bmatrix} L'_n & 2(F\otimes G)\tilde{D}_{n-1} \end{bmatrix}$$
$$= \begin{bmatrix} I_{n(n+1)/2} & 2L_n(F\otimes G)\tilde{D}_{n-1} \\ 2\tilde{D}'_{n-1}(F'\otimes G')L'_n & 4\tilde{D}'_{n-1}(F'\otimes G')(F\otimes G)\tilde{D}_{n-1} \end{bmatrix}$$
$$=: \begin{bmatrix} \mathcal{A} & \mathcal{B} \\ \mathcal{C} & \mathcal{D} \end{bmatrix}$$
(52)

we can compute the squared determinant of Z as

$$\begin{aligned} |Z'Z| &= |\mathcal{A}| |\mathcal{D} - \mathcal{C}\mathcal{A}^{-1}\mathcal{B}| \\ &= |\mathcal{D} - \mathcal{C}\mathcal{B}| \\ &= |4(\tilde{D}'_{n-1}(F' \otimes G')(I_{n^2} - L'_n L_n)(F \otimes G)\tilde{D}_{n-1})| \\ &= 2^{(n-1)(n-2)} |\tilde{D}'_{n-1}(F' \otimes G')(I_{n^2} - L'_n L_n)(F \otimes G)\tilde{D}_{n-1}|. \end{aligned}$$
(53)

Note that the exponent of 2 is smaller than in the Case of Inoue and Kilian [2018]. The reason for that is that because J_Q has less columns in our case Z has so too. Therefore the dimension of Z'Z is (n-1)(n-2)/2 here instead of n(n-1)/2. Now we define $F^* = (I_{n-1}-S^*)^{-1}$ and $\tilde{D}_{n-1}^+ = (\tilde{D}'_{n-1}\tilde{D}_{n-1})^{-1}\tilde{D}'_{n-1}$. Using a result from Magnus [1988] we can write

$$\tilde{D}'_{n-1}(F'\otimes G')(I_{n^2} - L'_nL_n)(F\otimes G)\tilde{D}_{n-1}) = \tilde{D}'_{n-1}((F^*)'\otimes (F^*)')(K'\otimes K')(I_n\otimes C')(I_{n^2} - L'_nL_n)(I_n\otimes C)(K\otimes K)(F^*\otimes F^*)\tilde{D}_{n-1} = \mathcal{M}'\tilde{D}'_{n-1}(K'\otimes K')(I_n\otimes C')(I_{n^2} - L'_nL_n)(I_n\otimes C)(K\otimes K)\tilde{D}_{n-1}\mathcal{M},$$
(54)

where $\mathcal{M} = (\tilde{D}_{n-1}^+(F^* \otimes F^*)\tilde{D}_{n-1})$. Again, using Magnus [1988] we obtain

$$|\mathcal{M}| = |F^*|^{n-2} = |\frac{1}{2}I_{n-1} + Q^*|^{n-2} = 2^{-(n-2)(n-1)}|I_{n-1} + Q^*|^{n-2}.$$
(55)

In the case of sign restrictions without exclusion restrictions, Inoue and Kilian [2018] are able to simplify (54) further by decompositioning $I_{n^2} - L'_n L_n = R'R$ where R is an $n(n-1) \times n$ matrix that satisfies $RR' = I_{n(n-1)/2}$. Further, they are able to show that

$$R(I_n \otimes C)\tilde{D}_n \tag{56}$$

is lower triangular and therefore they can simplify its determinant to a product of diagonal elements of C. Because

$$R(I_n \otimes C)(K \otimes K)\tilde{D}_{n-1} \tag{57}$$

is no square matrix, however, we cannot proceed accordingly and do not simplify our expression further. Using (46), (47), (51), (53), (54) and (55) we obtain

$$\begin{aligned} J_{1} &= |C|^{np} |(Q' \otimes I_{n}) L'_{n} (I_{n} \otimes C) J_{Q}| \\ &= |C|^{np} |(Q' \otimes I_{n}) Z| \\ &= |C|^{np} |Z| \\ &= |C|^{np} 2^{(n-1)(n-2)/2} |\tilde{D}'_{n-1}(F' \otimes G') (I_{n^{2}} - L'_{n} L_{n}) (F \otimes G) \tilde{D}_{n-1}|^{1/2} \\ &= |C|^{np} 2^{(n-1)(n-2)/2} |\mathcal{M}| |\tilde{D}'_{n-1}(K' \otimes K'C') (I_{n^{2}} - L'_{n} L_{n}) (K \otimes CK) \tilde{D}_{n-1}|^{1/2} \\ &= 2^{-(n-1)(n-2)/2} |C|^{np} |I_{n-1} + Q^{*}|^{n-2} |\tilde{D}'_{n-1}(K' \otimes K'C') (I_{n^{2}} - L'_{n} L_{n}) (K \otimes CK) \tilde{D}_{n-1}|^{1/2}. \end{aligned}$$

$$(58)$$

For

$$J3 = \frac{\partial vech(\Sigma)}{\partial vech(C)'} \tag{59}$$

we can use the result from Inoue and Kilian [2018] that

$$|J_3| = 2^n \prod_{i=1}^n c_{ii}^{n-i+1},\tag{60}$$

where c_{ii} is the *i*th diagonal element of C. Inoue and Kilian [2013] proved that

$$\left|\frac{\partial vec(\Phi)}{\partial vec(B)'}\right| = 1,\tag{61}$$

we can re-write (32) as

$$f(\hat{\Theta}) = |J_1|^{-1} |J_3| f(B|\Sigma) f(\Sigma) f(s^*)$$

= $2^{(n^2 - n + 2)/2} |C|^{-np} |I_{n-1} + Q^*|^{-(n-2)} \prod_{i=1}^n c_{ii}^{n-i+1}$
 $|\tilde{D}'_{n-1}(K' \otimes K'C') (I_{n^2} - L'_n L_n) (K \otimes CK) \tilde{D}_{n-1}|^{-1/2} f(B|\Sigma) f(\Sigma) f(s^*).$

TU Bibliothek, Die approbierte gedruckte Originalversion dieser Diplomarbeit ist an der TU Wien Bibliothek verfügbar wien Nour knowledge hub The approved original version of this thesis is available in print at TU Wien Bibliothek.

B Replication of the Kilian and Inoue (2013) Oil Model

Since I could not rely on any pre-built packages or code to estimate and identify my SVAR model, I wanted to make my own coding results comparable to a benchmark model to detect potential programming or conceptional errors. Slightly modifying my code (i.e. reducing dimensions, getting rid of the part involving sub-rotations, changing some identifying assumptions and substituting my own formula for the IRF posterior density for the one provided by Inoue and Kilian [2018]) allowed me to replicate the oil model of Inoue and Kilian [2013]. This model has the advantages that a) its identifying assumptions are similar to the ones used in my model, b) it relies on the same strategy for inference (i.e. reporting the modal IRF) and c) there is Matlab code online available, which has been corrected for various errors/shortcomings that have been detected in the last years. Figure 20 was generated using this original Matlab code. Figure 21, on the other hand, plots the structural IRFs that were computed using my own code.

Taking into consideration, that due to different random number generators and seeds the results cannot be expected to be exactly the same, the two Figures can be interpreted as being roughly identical. Since the resulting posterior distribution is relatively flat in some regions, in particular small differences in the reported modal model should not be put on the gold scale. Even though not every part of my code could be cross-checked in this way (especially the derivation of the posterior distribution in the case of sub-rotations), this result should hopefully strengthen confidence in the technical soundness of this paper.



Figure 20: Structural Impulse Response Functions of the Inoue and Kilian [2013] Oil Model. *Notes:* This plot is produced using the original Matlab Code provided together with the Kilian and Lütkepohl [2018] textbook. Since the code has been subject to some revisions, the plot slightly deviates from the Inoue and Kilian [2013] paper but resembles Figure 13.7 in Kilian and Lütkepohl [2018]. This code can be downloaded from Lutz Kilian's website: https://sites.google.com/site/lkilian2019/textbook/code?authuser=0.



Figure 21: Replication of the Inoue and Kilian [2013] Oil Model using my own R code. *Notes:* This figure was produced using the original dataset of Inoue and Kilian [2013] and a modified version of my own R code that was originally programmed to analyse the U.S. natural gas market. Since the exact result depends on the random number generator, the seed and the exact number of iterations, some minor differences to the original Figure are to be expected.

C Robustness Checks

C.1 Results for Subsamples



Figure 22: Structural IRFs for the early sample 1993M11-2005M04. *Notes:* Dashed lines mark IRFs corresponding to modal model of late sample 2005M05-2020M02 as in Figure 23.



Figure 23: Structural IRFs for the late sample 2005M05-2020M02. *Notes:* Dashed lines mark IRFs corresponding to modal model of late sample 1993M11-2005M05 as in Figure 22.





Figure 24: Forecast error variance decomposition for the early sample 1993M11-2005M04.







C.2 Different Elasticity Bounds

Figure 26: Structural IRFs for the full sample for robustness check.

Notes: Here we restricted the immediate response of industrial production to a gas demand shock to lie in the interval (-b, 0), where b is one third of the variance of the reduced form residual of industrial produciton. The black line is the modal model which maximizes the posterior density of the structural IRFs. Colored lines are the Bayesian 68% highest density region. 500 reduced form draws with 10,000 draws of Q^* for each reduced form draw were conducted.



Figure 27: Forecast error variance decomposition for the full sample for robustness check. Notes: Here we restricted the immediate response of industrial production to a gas demand shock to lie in the interval (-b, 0), where b is one third of the variance of the reduced form residual of industrial production. Forecast error variance decomposition corresponds to the impulse response function with maximal posterior density given in Figure 26.