EUROPEAN CENTRAL BANK

EUROSYSTEM

WORKING PAPER SERIES

NO 1432 / MAY 2012

MACROECONOMIC SHOCKS IN AN OIL MARKET VAR

by Marko Melolinna

In 2012 all ECB publications feature a motif taken from the €50 banknote.

NOTE: This Working Paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

Acknowledgements
I thank the participants at the Helsinki Center of Economic Research Econometrics Seminar of May 2011, and an anonymous referee for helpful comments.

Marko Melolinna
at European Central Bank, Kaiserstrasse 29, D-60311 Frankfurt am Main, Germany; e-mail: marko.melolinna@ecb.europa.eu

Address
Kaiserstrasse 29, 60311 Frankfurt am Main, Germany

Postal address
Postfach 16 03 19, 60066 Frankfurt am Main, Germany

Telephone
+49 69 1344 0

Internet
http://www.ecb.europa.eu

Fax
+49 69 1344 6000

ISSN 1725-2806 (online)

Abstract

This paper studies oil market and other macroeconomic shocks in a structural vector autoregression with sign restrictions. It introduces a new indicator for oil demand, and uniquely, performs a sign restriction set-up with a penalty function approach in an oil market vector autoregression. The model also allows for macroeconomic shocks in the US. The results underline the importance of the source of an oil shock for its macroeconomic consequences. Oil supply shocks have been less relevant in driving real oil prices, and had less of an effect on US inflation than demand shocks. Overall, the effects of oil shocks on US real activity have been relatively limited, as also highlighted by a counterfactual experiment of recent oil market developments.

JEL No.: C01, C32, E32

Keywords: oil demand shocks, oil supply shocks, business cycle, Bayesian econometrics

Non-technical summary

In recent literature on the effects of oil shocks on the macroeconomy, it has been acknowledged that it is not sufficient to study merely the effects of oil price shocks in isolation of other shocks affecting the economy. Therefore, models that can take into account the interdependencies of various key macroeconomic variables have gained in popularity. One typical class of these kinds of models offering a simple, yet effective way of studying the shocks is structural vector autoregression (SVAR) models. Some recent studies have used traditional recursive structures to identify oil demand and supply shocks, while many studies have also used sign restrictions. The latter typically means setting sign restrictions on some of the impulse responses based on economic theory, and then simulating the model with different simulation methods.

In the oil shock literature, the definition of supply and demand shocks is not necessarily clear-cut, and there have been a variety of methods used to elicit the effects of the different oil shocks. The weakness of VARs is that they are not structural models grounded precisely on macroeconomic theory, so it is difficult to prioritise between different methodologies. However, it does seem obvious that the more one can move towards identifying the typical shocks usually found in traditional business-cycle VARs, the more realistic one can hope the answers to be. Furthermore, defining an indicator for the demand pressures in the oil markets is a challenge, and something may be gained by exploring new kinds of demand indicators.

The current paper attempts to contribute to the ongoing discussion on the macroeconomic effects of oil shocks in the following ways. First, it introduces a new six-variable oil market VAR, including a new indicator for measuring oil demand in traditional oil market VARs. Second, the model takes into account demand and supply as well as monetary policy shocks in the US economy, thus allowing for a more realistic setting to study the effects of oil market shocks. Third, and uniquely, the paper studies a more comprehensive set of identification restrictions for the different oil shocks than has previously been conventional in the literature. Namely, the current paper uses a so-called penalty function approach to impose sign restrictions in the model. Finally, the paper introduces a counterfactual experiment to illustrate how the model can be used to draw conclusions on recent developments in oil supply and prices.

The results of the paper largely confirm those of recent literature, but with some interesting exceptions and additions. My results suggest that the US real economic activity responds differently to different oil shocks, namely, positive oil demand shocks have been associated with a positive effect on US GDP, while oil supply and expectational shocks have had a negative effect. Overall, oil shocks have not been very significant drivers of economic activity in the US, which has been more dependent on domestic macroeconomic shocks. In addition, the counterfactual experiment suggests that an increase in oil supply would not have been effective in avoiding the recession of the late 2000's in the US. Hence, there is little support in my results for the claim made in previous literature that monetary policy responses to oil shocks have caused large macroeconomic fluctuations in the US, or for the assertion that oil shocks were a major reason for the recession in the US in the late 2000's.

In line with most of the recent literature, I also find the response of the oil price to be dependent on the kind of shock that hits the oil market. Demand shocks have had larger and more persistent effects on the real oil price than supply shocks, although based on the results in the current paper, the relevance of speculation as an important driver of oil prices cannot be ruled out either. On the other hand, general macroeconomic shocks in the US have had a less persistent effect on the oil price. This reflects not only the fact that domestic shocks in the US economy aren't necessarily large enough to move the global oil market, but also the fact that forces moving the oil price can usually be traced back to the fundamentals of demand and supply in the oil market.

The objective of the current study is not to provide definitive answers on the macroeconomic effects of oil shocks, or on how to measure oil demand. Rather, it offers a slightly different way of studying these shocks and demand pressures than has been used in the recent literature, and largely confirms the results of the emerging consensus. It also underlines the fact that different modelling and sign restriction strategies can have a significant influence on the inference of the results. However, all in all, a deeper understanding of the nature and definition of shocks hitting the oil market, and their policy implications for oil-importing economies is required. Clearly, as suggested by most recent studies on the macroeconomic effects oil shocks, modelling oil price shocks without taking into account the source of the shock is not a viable avenue.

1 Introduction

In recent literature on the effects of oil shocks on the macroeconomy, it has been acknowledged that it is not sufficient to study merely the effects of oil price shocks in isolation of other shocks affecting the economy. Therefore, models that can take into account the interdependencies of various key macroeconomic variables have gained in popularity. One typical class of these kinds of models offering a simple, yet effective way of studying the shocks is structural vector autoregression (SVAR) models. One of the seminal papers of recent literature of this branch is Kilian (2009), which introduces a recursively structured three-variable oil market SVAR. The paper identifies the model by assuming that oil supply does not respond to innovations in oil demand during the same month, and by inserting the identified shocks into separate models explaining US GDP growth and inflation, concludes that these effects depend crucially on what kind of a shock we are dealing with. In particular, demand shocks that increase oil prices can have positive effects on US GDP. Furthermore, the paper concludes that empirical monetary VARs linking the response of monetary policy to oil shocks, like the one introduced by Bernanke et. al. (1997) are fundamentally misspecified, because they assume the same response to oil price innovations regardless of the composition of the shock. Other authors, like Hamilton and Herrera (2004), Kilian (2010) and Kilian and Lewis (2010) have also challenged the findings of Bernanke et. al. (1997). Hence, the claim in Bernanke et. al. (1997) that considerable aggregate macroeconomic fluctuations have been caused by monetary policy responses to oil shocks, is open to debate.

Instead of applying recursive structures to identify the VAR models, many studies have also used sign restrictions. This typically means imposing sign restrictions on some of the impulse responses based on economic theory, and then simulating the model with, for example, Monte Carlo simulation methods, as suggested by Uhlig (2005) and Mountford and Uhlig (2009). For example, Kilian and Murphy (2009) apply sign restrictions to the recursive basic oil market model introduced by Kilian (2009), and hence try to overcome the slightly implausible recursive structure of the latter paper. By using sign restrictions as well as some restrictions on short-term supply elasticity, the authors show that the results of Kilian (2009) are mostly valid, and that the real price of oil has been largely driven by demand rather than supply shocks. Kilian and Murphy (2011) further develop the basic oil market model to take into account speculation (measured with data on oil stocks) in a sign-restriction framework, but still conclude that the fluctuations seen in the real price of oil during the past 10 years or so were caused mainly by demand rather than supply or speculation shocks. Lippi and Nobili (2009) find that the effects of oil shocks on the main macroeconomic variables depend on the source of shock, with a positive oil demand shock being associated with an increase in economic activity in the US.

Some authors have used more advanced VAR techniques in eliciting the effects of different kinds of oil shocks. For example, Aastveit (2009) uses a data-rich environment in a recursively identified factor-augmented VAR (FAVAR) approach to conclude that oil demand shocks are more important than supply shocks as a driving force behind several macroeconomic variables. Furthermore, he also finds that monetary policy in the US reacts differently to different kinds of oil shocks. Baumeister and Peersman (2011), using a time-varying VAR framework, find a large role for oil demand shocks in real oil price variability, and this role has increased over time.

However, despite the abundance of studies in the field, none of the models mentioned above is without its flaws. As the above discussion indicates, the definition of supply and demand shocks is not necessarily clear-cut, and there have been a variety of methods used to elicit the effects of the different oil shocks. Obviously, the weakness of VARs is that they are not structural models grounded precisely on macroeconomic theory, so it is difficult to prioritise between different methodologies. However, it does seem obvious that the more one can move towards identifying the typical shocks usually found in traditional business-cycle VARs, the more realistic one can hope the answers to be. Also, one additional weakness of some of the models used in the previous literature is the sign restriction methodology, which, as shown by Fry and Pagan (2007), suffers from certain theoretical drawbacks. Hence, it would be useful to study oil market VARs with different sign restriction strategies. Furthermore, defining an indicator for the demand pressures in the oil market is a challenge, which seems to have been met in the recent literature by the indicator introduced by Kilian (2009). Nevertheless, something may be gained by exploring other kinds of demand indicators.

The current paper attempts to contribute to the ongoing discussion on the macroeconomic effects of oil shocks in the following ways. First, it introduces a new six-variable oil market VAR, including a new indicator for measuring oil demand in traditional oil market VARs. Second, the model takes into account demand/supply and monetary policy shocks in the US economy, thus allowing for a more realistic setting to study the effects of oil market shocks. Third, and uniquely, the paper studies a more comprehensive set of identification restrictions for the different oil shocks than has previously been conventional in the literature. Namely, the current paper uses a so-called penalty function approach, originally introduced by Uhlig (2005), to impose sign restrictions in the model. Finally, the paper introduces a counterfactual experiment to illustrate how the model can be used to draw conclusions on recent developments in oil supply and prices.

The results of the paper largely confirm those of recent literature, but with some interesting exceptions and additions. As has been concluded by many previous studies, I also find that the source of the oil shock matters for its macroeconomic consequences. In particular, supply shocks tend to be more deflationary than demand shocks, and oil demand shocks tend to have a more persistent effect on the oil price. Overall, oil shocks have not been very significant drivers of economic activity in the US, which has been more dependent on domestic macroeconomic shocks. In addition, the counterfactual experiment suggests that an increase in oil supply would not have been effective in avoiding the recession of the late 2000's in the US.

The paper is organised as follows. Section 2 introduces the model, data as well as the details of the oil demand indicator. Section 3 deals with the sign restriction strategy, and section 4 with the results. Section 5 describes the counterfactual experiment. Section 6 concludes.

2 The model

As the previous discussion implies, oil shocks have been included in VAR models in a variety of ways in the previous literature. From a theoretical perspective, none of the methods is *ex ante* better than another one. The number of variables used in the models varies from three in the simple model introduced by Kilian (2009) to many dozens in the FAVAR model used, for example, by Aastveit (2009).

The current study attempts to strike a balance between the simplicity of the model and the applicability of it for studying not only oil market shocks, but also the interaction between oil market variables and other key macroeconomic variables. The idea is to illustrate how such a simple model can be used for inference by analysing impulse response functions and also a simple counterfactual experiment. The model used in the analysis is a six-variable SVAR model concentrating on the US economy, with the aim of including key shocks affecting the US business cycle as well as the global oil market. The variables included in the model are global oil production, an oil demand indicator, the real price of oil, a US GDP indicator, US inflation and US monetary policy rate. The model is similar in structure to monetary oil market models used in the recent literature (see, for example, Peersman and Van Robays (2009)).

The methodology of the paper is based on a simple reduced-form VAR model of the following type:

$$Y_t = \sum_{i=1}^{L} B_i Y_{t-i} + u_t \tag{1}$$

where Y_t is an $Nx1$ vector of endogenous variables (so in my model $N = 6$), L is the lag length of the VAR, B_i are the NxN coefficient matrices, and u_t is a vector of white noise reduced-form innovations with a variance-covariance matrix Σ_u.

2.1 Data

The variables included in the model are global oil production, an oil demand indicator (see Section 2.2 for more details on the demand indicator), the real price of oil (West Texas Intermediate brand), a monthly US GDP indicator (measured by the Chicago Fed National Activity Index), US inflation and US monetary policy rate (proxied by 3-month interbank market rate). All the data are in stationary form: inflation and GDP are measured by year-on-year change, the supply and demand indicators are in deviation from a long-term trend, the oil price is in (log) level, and the interest rate is in level. The sample for the model is 1974M1 to 2010M12, and there are 12 lags in the model[1]. The model also includes a constant and 11 seasonal dummies[2], which are suppressed for simplicity in equation (1).

2.2 Coincident oil demand indicator

Using the oil demand indicator introduced by Kilian (2009) has become common practice in recent oil market VAR models. This indicator measures oil demand pressures by combining global bulk dry shipping freight rates to form a single index. However, this indicator (henceforth K09 indicator), whilst offering a good proxy for general demand pressures in the commodities markets on a monthly frequency, suffers from certain drawbacks. The K09 indicator is heavily dependant not only on demand pressures, but also on supply of cargo ships, which might be expected to lag the economic cycle somewhat. Furthermore, it is not clear why dry shipping freight rates should react to some of the shocks in the way typically imposed on the sign restriction VAR literature (for more details, see Section 3.1 below). For example, it is not unambiguous why a shock of a precautionary increase in oil demand would lead to a decrease in these freight rates. If anything, the increase in oil demand might lead to upward pressures in all kinds of shipping freight rates.

Consequently, I introduce a new alternative oil demand indicator. This monthly coincident demand indicator is based on variables that can be expected to reveal relevant information about oil market demand pressures. The coincident demand indicator is based on monthly industrial production and petrol consumption data, i.e., economic activities that directly and instantly affect the demand for oil. I construct the indicator by combining indicators that can be expected to affect the demand, in a state-space framework into a single demand indicator a la Stock and Watson (1991) (for more technical details, see Appendix A). Some of these indicators, especially for emerging markets, are not available for the entire sample history used in my VAR model. However, the state-space framework allows for missing values to be filled in by the recursive Kalman filter estimation technique.

The indicator used in the model is as parsimonious as possible and has only three variables; OECD industrial production volume (sourced from the OECD), US personal consumption expenditures on oil (the US Bureau of Economic Analysis), and emerging markets industrial production volume (CPB Netherlands Bureau for Economic Policy Analysis). These subcomponents are depicted (as deviations from long-term trends) in Figure 1. The figure shows that while the direction of the subcomponents has not always been similar, the demand indicator captures the general direction, and especially the importance of emerging market demand in the early part of

[1] Traditional information criteria favour a shorter lag structure. However, as pointed out by Hamilton and Herrera (2004), it is important to have sufficiently long lag structures in oil market VARs. In any case, the results are qualitatively robust to both shorter and longer lag structures.

[2] Even though none of the variables is expected to necessarily show a strong seasonal pattern, as is somtimes customary in oil market models of this type (see, for example, Kilian and Murphy (2011)), seasonal dummies are included. However, the results are very similar without the dummies.

the 2000's quite well.

The coincident demand indicator, together with the K09 indicator, is illustrated in Figure 2, as deviations from the long-term trend. As can be the seen in the graph, the differences between the two indicators are normally not large. However, during the 2000's, the coincident demand indicator shows a more steady increase, and then the drop in the latter half of the decade is more pronounced in the K09 indicator.

The choice of the demand indicator is, of course, arbitrary, and other indicators might also be suggested. However, the coincident demand indicator has advantages over the K09 indicator. First, it attempts to capture variables of the global economy that are directly affected by oil, but also carry a macroeconomic interpretation. Second, the coincident demand indicator has a higher correlation with measures of global GDP growth and business cycles than the K09 indicator; over the sample horizon of the current study, the correlation coefficient of the coincident demand indicator with global GDP growth (on a quarterly basis) is 0.42 and with the global business cycle 0.61, whereas the corresponding correlations for the K09 indicator are 0.09 and 0.33, respectively[3]. While not proving any definite superiority of the coincident demand indicator, these facts at least suggest a place for it as a viable alternative[4].

While the choice of the demand indicator is unlikely to cause dramatic changes for inference in large VAR models, it can still have a substantial effect in some models for some impulse responses. As an example, this is highlighted in Figure 3 for the impulse response of oil production to an oil demand shock in the simple 3-variable structural oil market model (which is based on a Choleski identification scheme) introduced by Kilian (2009). As can be seen from the chart, after the initial impact period (when, by construction, oil supply does not respond to an oil demand shock), oil supply is much more responsive to the shock, when oil demand is measured by the coincident demand indicator than when it is measured by the K09 indicator. In fact, with the K09 indicator, the response isn't statistically significant apart from one period. While oil production can typically be expected to respond sluggishly to oil demand shocks, the positive response with the coincident demand indicator is perhaps more intuitive.

[3] The quarterly global GDP time series is derived by linking ECB data (available from 1981 onwards) with GDP growth rates in OECD countries for the pre-1981 period. The business cycle is computed by calculating a trend with a Hodrick-Prescott filter.

[4] One alternative to compare different oil demand indicators would be to see how they react in situations where there are demand pressures in an oil market producing oil near full capacity. Unfortunately, such capacity measures are not available for the length of the sample.

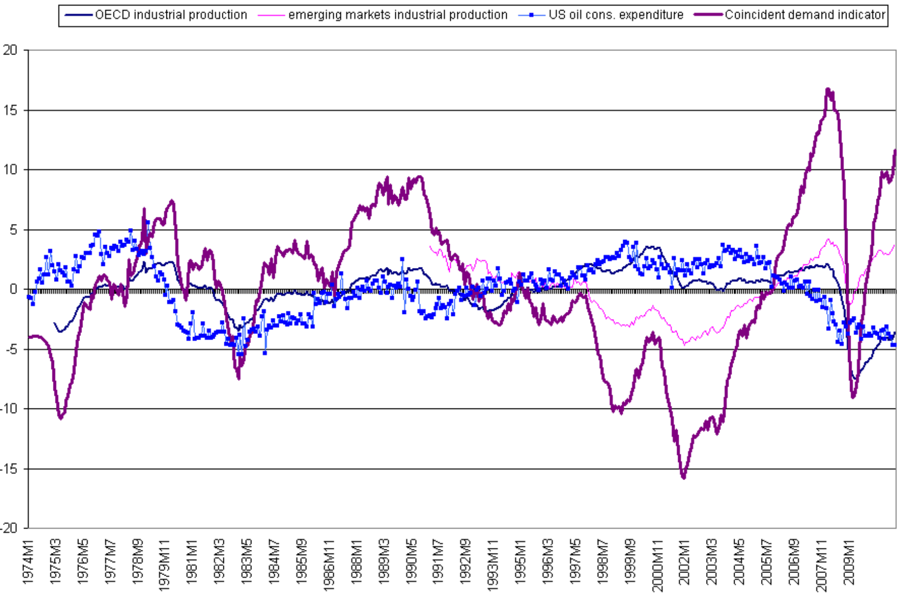

Figure 1: Subcomponents of the oil demand indicator

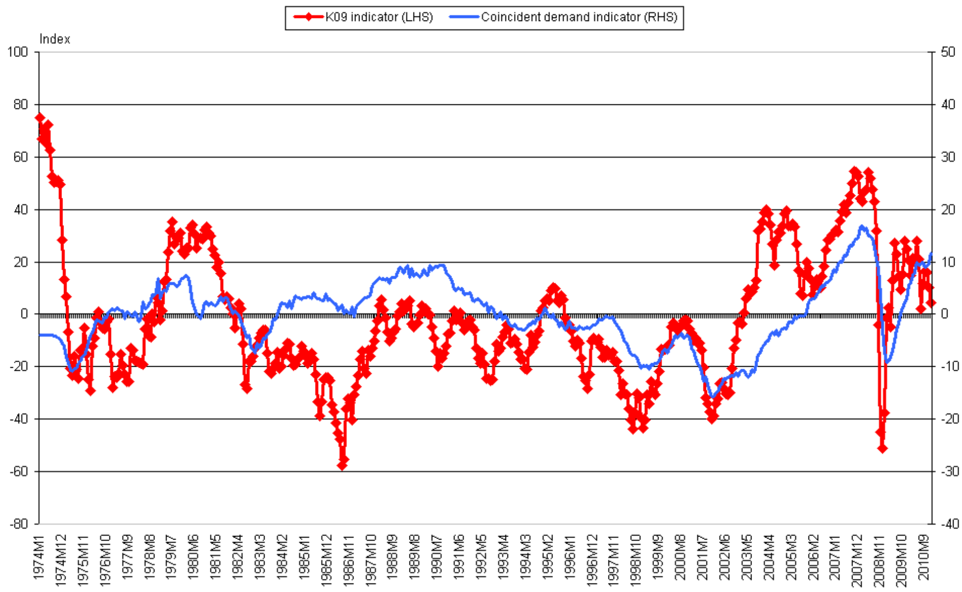

Figure 2: Oil demand indicators

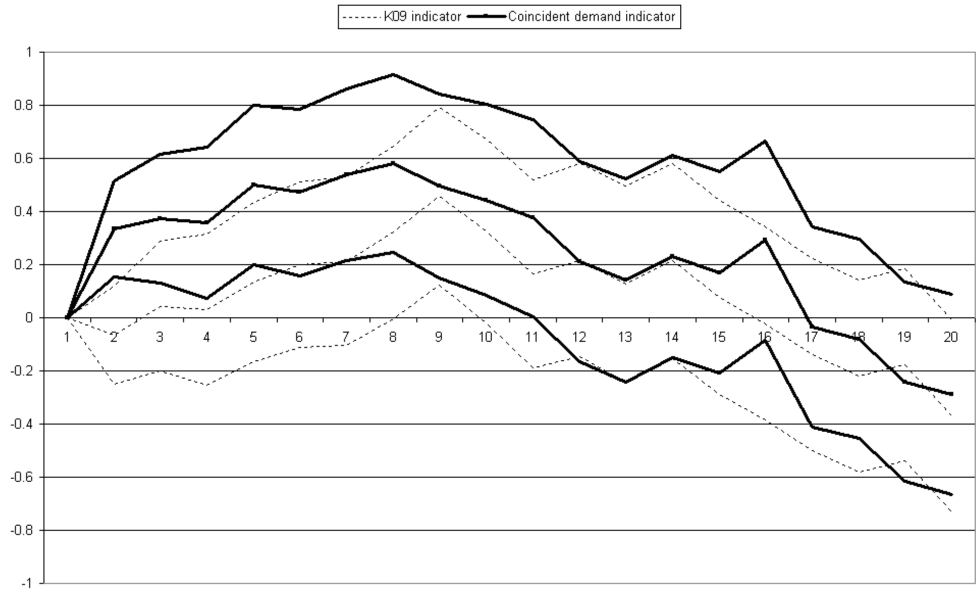

Figure 3: Impulse responses (in months) of oil production to an oil demand shock in Kilian (2009) oil model. Note: 16th and 84th percentile confidence intervals.

3 Shock identification

3.1 Sign restriction motivation

In reality, oil prices are constantly affected by different kinds of shocks of different sizes, which are priced in the markets intra-day. Due to the difficulty of modelling all the intra-day shocks, it is conventional to concentrate on a few lower-frequency shocks which can be easily defined. As in any market, these shocks are usually demand and supply shocks and possibly some form of expectational shocks. Following Kilian (2009), in most recent oil market VARs there are usually three different kinds of shocks than can affect the oil market. First, there are oil supply shocks, which are shocks affecting the current physical availability of oil. Second, there are aggregate demand shocks, which are defined as shocks to the current demand of oil driven by fluctuations in the global business cycle. Finally, there are precautionary demand shocks, which are shifts in precautionary demand for oil driven by expectational shocks to the demand and supply balance in the oil market.

My approach for identifying the oil shocks follows these definitions. However, unlike most other studies, I also allow for selected macroeconomic shocks (namely, a US demand/supply shock and a monetary policy shock) in the model[5]. Hence, whereas the traditional oil market models

[5]In this respect, my model is similar in spirit to those of Aastveit (2009), which allows for a monetary policy shock in addition to the three oil market shocks, and of Lippi and Nobili (2009), which allows for demand and

typically allow for business cycle fluctuations to affect the oil market through an aggregate demand shock, I will include selected macroeconomic shocks in my model explicitly. This, I believe, will offer a more realistic view on the interactions between different kinds of non-oil shocks typically hitting a macroeconomy (in this case, the US), and the oil market. This feature of my model also follows the spirit of the methodology introduced by Mountford and Uhlig (2009) to study the effects of fiscal policy shocks in a business-cycle VAR.

Unlike, for example Kilian (2009), my modelling strategy also makes an effort to move away from traditional, purely recursive ordering schemes. In this respect, my sign restriction modelling strategy is closer to oil market model identification strategies introduced by Baumeister and Peersman (2009, 2011), Peersman and Van Robays (2009) and Lippi and Nobili (2009). Although some degree of ordering is required, restrictive (and potentially incorrect) ordering of the different kinds of shocks in the oil market is avoided.

The sign restrictions assigned to the shocks in the model are presented in Table 1. Just like in Mountford and Uhlig (2009), there are two kinds of shocks in the model. In my model, these are oil market shocks and general macroeconomic shocks for the US economy. The sign restrictions imposed on the oil shocks follow the conventions of the previous literature, and hence require little justification. Negative oil supply shocks can be expected lift the oil price and if anything, lower the oil demand indicator. Both types of demand shocks will naturally lead to an increase in the oil price and (at least eventually) oil production. The pure oil demand shock is driven by an increase in the demand indicator, while the precautionary shock will lead to a decrease in the demand indicator (or, to be more precise, in economic activities requiring crude oil), as demand is directed away from its economic use to a precautionary use[6]. Note that the oil market shocks are all defined as shocks that lift the oil price (i.e., either a cut in production or an increase in demand) and that the model is agnostic about the effects of the oil shocks on US GDP.

The general macroeconomic shocks also follow previous conventions in the literature. US demand shock (which can be seen as a positive business cycle shock) lifts both activity and prices in the US, while a supply shock (which can be interpreted as a positive productivity shock) lifts activity, but lowers prices. The monetary policy shock (a monetary policy rate hike) has a negative effect on prices and activity. Note that the model is agnostic about the effects of the US macroeconomic shocks on the oil market variables, since even though shocks in an economy the size of the US can be expected to have an effect on these variables, it cannot be presumed.

All the shocks in the model are one standard deviation changes in the restricted variables. All the sign restrictions are required to apply for 6 months (i.e., $K = 5$ in equation (5)), which is

supply shocks both in the oil market and the US economy.

[6] In the benchmark model, I am not interested in modelling speculation and/or inventory behaviour explicitly, so no inference on these factors can be drawn. Some alternatives for this are studied below. For studies addressing these issues more explicitly, see, for example, Kilian and Murphy (2011), or Lombardi and Van Robays (2011).

standard in the literature[7]. This time period can be expected to be long enough for the shocks to have the desired effects on the restricted variables.

	Oil production	Oil demand indicator	Real oil price	US GDP	US inflation	US short rate
Oil market shocks						
Oil supply shock	-	-	+		+	
Oil demand shock	+	+	+		+	
Oil precautionary shock	+	-	+		+	
Non-oil market shocks						
US demand shock				+	+	
US supply shock				+	-	
Monetary policy shock				-	-	+

Table 1: VAR model sign restrictions

In addition to the sign restrictions, I also restrict the supply elasticity after both of the oil demand shocks. In my model, as well as in the previous literature, supply elasticity is defined as the ratio of the change in oil supply divided by the change in the real price of oil for an oil demand shock. Hence, as the variables are in logs, the supply elasticity can be defined as the structural impact response (i.e., the impact period coefficient in the structural model) of oil supply divided by the impact response of the real price of oil for an oil demand shock. According to the previous literature, this elasticity can be expected to be small, as oil production typically responds slowly to changes in price. According to Kilian and Murphy (2009), this elasticity could be somewhere between 0 and 0.025 (i.e., a one percent increase in the price would cause maximum a 0.025 percent increase in supply). However, since I want to remain agnostic about the possible magnitude of the elasticity, I only restrict[8] it to lie between 0 and 0.2.

3.2 Sign restriction strategies

Generally, following Uhlig (2005), there are two different types of sign restriction strategies used in VAR models; a pure sign-restrictions (PSR) approach and a penalty function approach[9]. Traditionally, in previous oil market models, sign restriction strategies have concentrated purely on the PSR approach. As has been widely pointed out in the literature (see, for example, Fry and Pagan (2007)), deriving median responses based on the accepted draws in the PSR approach is theoretically incorrect. This is because the accepted responses come from different models, due to the fact that coefficients of the model change for each draw. Furthermore, even though some solutions have been suggested for overcoming these problems in the PSR approach, my model

[7]However, the results presented below are qualitatively robust to different values of K, like $K = 11$, which is another popular option in the literature.

[8]In the benchmark model simulations detailed below, this restriction rejects about 75% of the draws.

[9]Appendices B and C describe these two methodologies in more detail.

is so large that searching for accepted draws with the PSR approach becomes very cumbersome in practice. This is due to the fact that even excluding the supply elasticity restrictions the benchmark model includes 114 sign restrictions altogether (the 19 sign restrictions in Table 1 over 6 months)[10]. In addition, it is important to notice that the sign restrictions do not uniquely identify all the shocks; in particular, the US demand shock could be observationally equivalent to any of the oil market shocks, since the sign restrictions do not preclude each other. Hence, in my benchmark model, I will concentrate on the penalty function approach. To my knowledge, this is a unique feature of my model, as this approach has not been applied to oil market VARs in the previous literature.

As regards carrying out the penalty function approach, and following Mountford and Uhlig (2009), I first identify the oil market shocks via a penalty function and sign restrictions, thus ascribing as much movement in the oil market variables to these shocks as possible. This is intuitive, since oil market dynamics can be expected to be mostly affected by shocks in oil demand and supply directly. The US macroeconomic shocks are then identified via sign restrictions as well as orthogonality restrictions to the oil market shocks. These orthogonality restrictions also allow the shocks to be uniquely identified, unlike in the PSR approach. This is because the penalty function procedure is somewhat reminiscent of a causal ordering (like a Choleski decomposition). If the penalty function was linear, it would be possible to use a Choleski decomposition to minimise the penalty function. Since the penalty function used here is, however, non-linear and involves the impulse responses for several periods, this is not possible, but the analogy may still be helpful to understand the procedure and the results (see Mountford and Uhlig (2009)).

As Fry and Pagan (2007) point out, the theoretical problems related to the median responses of the PSR approach are averted in the penalty function approach, as long as the criterion function is a reasonable one. However, it must be of course recognised that this is an additional restriction on the model, and more than just sign restriction information is used for inference. However, I maintain that the additional restrictions used are fairly loose and intuitive, and hence the sign restriction strategy is viable.

4 Results

4.1 Benchmark results

As mentioned above, the identification strategy used in this study is based on sign restrictions. As regards the PSR approach, apart from the theoretical problems detailed above, the identification tends to be relatively weak, as confidence bounds are large[11]. Hence, the PSR approach would not appear to be a constructive method in this kind of a model.

[10] Simulations suggest that with these sign restrictions, it will take around 100,000 draws to generate 1 accepted draw. Thus searching for a sufficient amount of accepted draws becomes an extremely time-consuming task.

[11] These results are available from the author on request.

The impulse responses for the penalty function approach (with 16th, 50th and 84th quantile confidence bounds) are presented in Appendix D. The impulse responses are based on 1,000 accepted draws. This sign restriction methodology provides some interesting results, and the responses not restricted for their signs appear intuitive. As regards the oil market shocks (ordered first in the model), an oil supply shock that cuts production, causes the oil price and US inflation to increase, as well as the US economy to contract in the short term. The macroeconomic effects in the US of a precautionary oil demand shock are fairly similar to those of the supply shock, except that the inflationary effect is more persistent, which also leads to a more positive reaction of the monetary policy interest rate. The oil demand shock, on the other hand, leads to a more pronounced effect on the real price of oil, suggesting that pure demand shocks have been more important in driving the oil price than supply shocks. Furthermore, the oil demand shock has a positive effect on US GDP in the short term, probably reflecting the favourable global economic situation related to positive oil demand shocks. The effect on US inflation and monetary policy is also more pronounced than in the other two oil market shocks.

As for the US macroeconomic shocks, the effects on the domestic economy are largely in line with the previous literature. The US demand shock has a fairly significant effect on the real oil price, but this effect is not as large or persistent as those caused by oil market shocks. The effects of the US monetary policy shock on the real price of oil are relatively limited. It is also worth noting that the US supply (i.e. productivity) shock leads to an increase in oil demand and price, even though these effects are not as large as for the US demand shock.

4.2 Robustness checks

The following two subsections present the results of robustness checks on the model. In particular, the issue of the robustness of the results to different specifications of the penalty function is examined. Furthermore, the results for a model including oil stocks, as suggested by Kilian and Murphy (2011), are also reported.

4.2.1 General robustness checks

As noted above, it must recognised that the penalty function imposes an additional restriction on the model, and hence its validity requires some justification. Here, similar justifications to those used by Uhlig (2005) in the case of monetary policy shocks can be used. First, the reward in the penalty function given to responses satisfying the sign restrictions is plausible, because I want to ensure that the responses caused by the shocks are economically meaningful, and not merely the product of some other, smaller, possibly random shocks hitting the economy. Second, because I want to impose sign restrictions, the penalty function should be asymmetric, punishing violations a lot more strongly than rewarding large and correct responses. Third, for the minimisation procedures to work, the penalty function needs to be continuous. Fourth, I want to punish even small violations (which is why e.g. a quadratic functional form is less appealing than the linear form), but also, I want to punish large deviations more than small ones (which is why e.g. a square root specification is less appealing). However, to examine the robustness of my results, I have also carried out the estimations in forms where the x in the penalty function is quadratic (i.e. x^2), square root ($\sqrt{|x|}$) and symmetric ($f(x) = x$ if $x > 0$). In all these robustness checks, the results stay qualitatively similar to the benchmark specification. Hence, even though the benchmark penalty function is to a certain extent *ad hoc*, the results are robust to changes in its form.

4.2.2 Model with oil stocks

In the spirit of Kilian and Murphy (2011), the model was also extended to include a seventh variable, namely oil stocks[12]. In this model, following Kilian and Murphy (2011), I define a speculative demand shock as a precautionary shock, which leads agents to accumulate crude oil stocks (for example, in expectation of a forthcoming supply disruption)[13]. This will lead to an increase in the oil price, but since the stock accumulation disrupts current oil availability, will also lead to an increase in oil production and/or a decrease in the oil demand indicator. The other shocks in the model follow the 6-variable benchmark model. It needs to be emphasised that the other two oil market shocks, the oil demand and oil supply shock, relate to shocks in *current* oil flows. The sign restrictions applied to the oil market shocks (the non-oil market shocks are suppressed for simplicity) are presented in Table 2. The model also follows other restrictions set out by Kilian and Murphy (2011), namely dynamic sign restrictions (although I am slightly more agnostic about the length of the shocks and consistently with the benchmark model, allow for $K = 5$), price elasticity of oil supply and price elasticity of oil demand[14].

The results of the model are largely very similar to those obtained by the benchmark model, which further confirms the robustness of the results to different specifications. In Appendix E, I only present the impulse responses of the oil market variables to oil market shocks. According to these results, the speculative demand shock has had a relatively persistent effect on the oil price. Furthermore, this shock leads to a more persistent increase in oil stocks. This is intuitive, and in line with the results of Kilian and Murphy (2011), as the precautionary motive prompts agents to hold oil stocks for an extended period of time.

[12] Following Hamilton (2009) and Kilian and Murphy (2011), global oil stocks are approximated by using OECD stock data from 1988 onwards, and by using US crude oil stocks as a scaling factor of OECD stocks between 1974 and 1987. All other data in the model follow the 6-variable benchmark model.

[13] Note that the speculative shock is agnostic about its cause, and even though the model does not include oil futures markets explicitly, arbitrage will ensure that futures price increases lead agents to buy inventories in the spot market.

[14] Kilian and Muprhy (2011) introduce the concept of oil demand elasticity in use (η_t), which is a time-varying measure taking into account the response of both current oil production and oil stocks to an oil supply shock. Following the authors, I impose the restriction of $-0.8 \leq \eta_t \leq 0$. For more details on how to derive the measure, see Kilian and Murphy (2011).

	Oil production	Oil demand indicator	Real oil price	Oil stocks
Oil market shocks				
Oil supply shock	-	-	+	
Oil demand shock	+	+	+	
Oil speculatuve shock	+	-	+	+

Table 2: 7-variable VAR model oil market sign restrictions

5 An example of a counterfactual analysis

To illustrate another fruitful way of using an oil market VAR model, which also includes macroeconomic variables, this section considers a counterfactual conditional forecasting example. For this experiment, the 6-variable VAR model described above is first estimated from 1974M1 to 2004M6. Around this time, oil demand and the real oil price started their long ascent, while oil production growth faded (see Figure 4). The counterfactual experiment examines the question of what the consequences would have been of oil supply growth continuing at levels seen just before the latter half of 2004. Specifically, the counterfactual experiment assumes that for 2004M7 to 2008M6, the year-on-year oil production growth would have stayed constant at 5%. Of course, it is improbable that capacity constraints and/or OPEC decision-making would have allowed production to stay at such levels during this period, but nevertheless, this experiment can give an insight into the potential effects of the rapid oil price increase of the latter half of the 2000's.

Technically, from 2004M7 onwards, the experiment is carried out by estimating the model parametres (called y_a henceforth) conditional on the *actual* oil production numbers, and then estimating the parametres (called y_{cf}) conditional on the *counterfactual* production numbers[15]. Then, the paths taken by selected variables of interest in the counterfactual scenario are subtracted from the paths in the actual scenario (i.e., $y_{cf} - y_a$). It might seem logical instead to compare *actual* model model parametres (in other words, not conditionally forecasted on actual oil production numbers) to the counterfactual experiment. However, this approach would mix all the other sources of fluctuation in the parametre values with the pure effects of oil production, and hence, would not help in eliciting the effects of the counterfactual experiment.

A potential caveat in the analysis is the dependence on the assumption that for the validity of the analysis, the coefficients of the model cannot have changed between 2004M7 and 2008M6. If they have, the paths would mix effects of the counterfactual analysis with changes in the behavioral relationship (Lenza et. al. 2010)). There is evidence that, even though the coefficients may have changed quite dramatically in the more distant past, the changes in the more recent past

[15] This strategy has been used recently, for example, in the literature studying the effects of non-standard monetary policy measures (see Lenza et. al. (2010)).

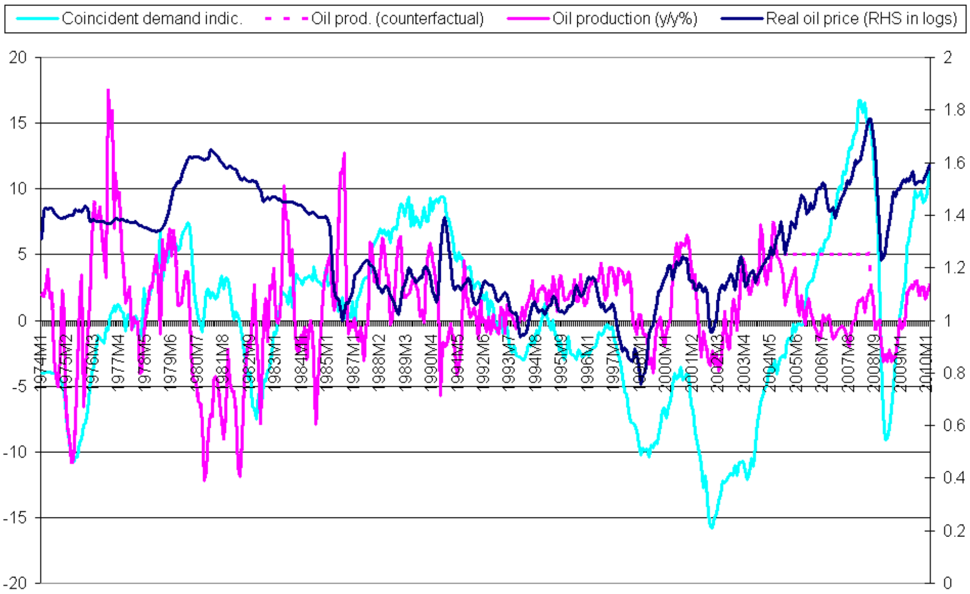

Figure 4: Counterfactual oil production assumption and relevant oil market variables

have been more modest (see Baumeister and Peersman (2011)). Furthermore, a traditional Chow breakpoint test strongly rejects a structural break in the middle of 2004. Hence, I consider this evidence enough to support the assumption and consequently the validity of the counterfactual experiment.

The results of the experiment (i.e. $y_{cf} - y_a$) for the US GDP, inflation and the real oil price are presented in Figure 5. The paths of the variables are intuitive; in the counterfactual scenario, the real oil price would have been nearly 10% lower by the beginning of 2008, while US GDP would have been slightly higher, and inflation initially lower than in the actual scenario. Based on this analysis, considering the subsequent steepness of the fall in US GDP after the recent recession set in, a lower oil price caused by higher supply would not have been enough to prevent the recession. Hence, this counterfactual scenario does not support the claim made, for example, by Hamilton (2009) that the increase in the oil price was a major cause of the US recession in the late 2000's, or at least that this effect could have been prevented by higher oil production. The result also further emphasises the role of demand as the main driver of the strong rise in the oil price in the 2000's.

6 Conclusions

This paper carries out an analysis of oil market shocks in a traditional business-cycle VAR. It makes three main contributions to the existing literature. First, it introduces a new six-variable

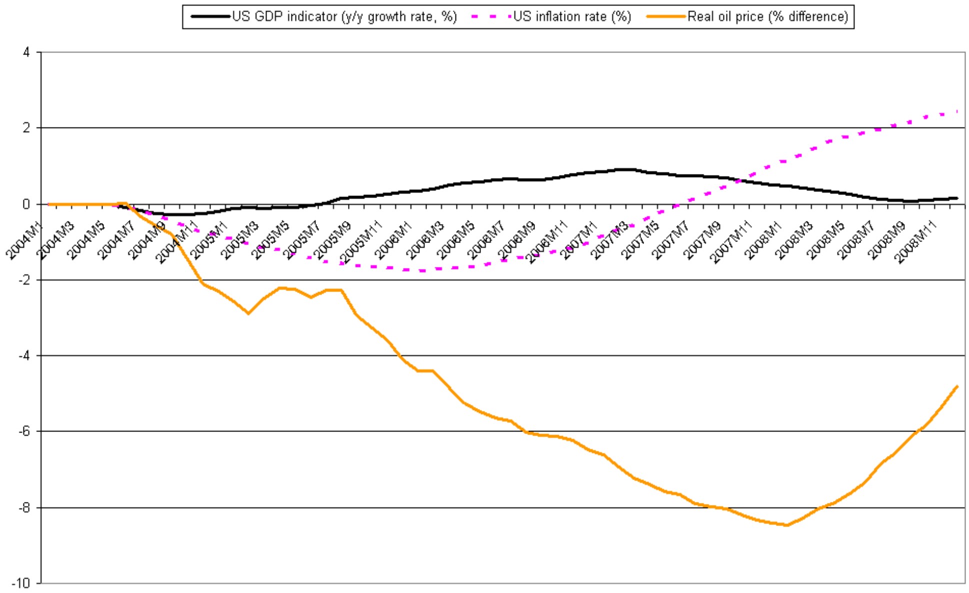

Figure 5: The difference between the selected VAR model variables under actual and counterfactual forecasting scenarios

business cycle oil market VAR, with a new indicator for measuring oil demand. The model takes into account macroeconomic shocks in the US economy, thus allowing for a more realistic setting to study the effects of oil market shocks. Second, the paper identifies the shocks in the model with a penalty function approach, instead of the pure sign restriction approach of previous studies. Finally, the paper introduces a counterfactual experiment to illustrate how the model can be used to draw conclusions on recent oil market developments.

The results of the analysis largely confirm those of previous studies, although not without exceptions. The source of the oil shock matters for its macroeconomic consequences. Oil supply shocks have tended to be less meaningful in moving inflation than demand shocks in the US economy over the past 35 years. Demand shocks have typically produced a more long-lasting effect on the US inflation and monetary policy. This also conforms well with the idea that monetary policy should react more to demand shocks that produce permanent upward pressure on prices, than to supply shocks, which cannot be mitigated by monetary policy.

My results suggest that the US real economic activity responds differently to different oil shocks, namely, positive oil demand shocks have been associated with a positive effect on US GDP, while oil supply and expectational shocks have had a negative effect. Overall, however, the response of real activity to oil shocks has been relatively muted compared to responses to the domestic macroeconomic shocks in the US. There is little support in my results for the claim in Bernanke et. al (1997) that monetary policy responses to oil shocks have caused large macroeconomic fluctuations in the US, or for the assertion in Hamilton (2009) that oil shocks were a major reason for the recession in the US in the late 2000's.

In line with most of the recent literature, I also find the response of the oil price to be dependent on the kind of shock that hits the oil market. Demand shocks have had larger and more persistent effects on the real oil price than supply shocks, although based on the results in the current paper, the relevance of speculation as an important driver of oil prices cannot be ruled out either. On the other hand, general macroeconomic shocks in the US have had a less persistent effect on the oil price. This reflects not only the fact that domestic shocks in the US economy aren't necessarily large enough to move the global oil market, but also the fact that forces moving the oil price can usually be traced back to the fundamentals of demand and supply in the oil market.

The objective of the current study is not to provide definitive answers on the macroeconomic effects of oil shocks, or on how to measure oil demand. Rather, it offers a slightly different way of studying these shocks and demand pressures than has been used in the recent literature, and largely confirms the results of the emerging consensus. It also underlines the fact that different modelling and sign restriction strategies can have a significant influence on the inference of the results. However, all in all, we still require a deeper understanding of the nature and definition of shocks hitting the oil market, and their policy implications for oil-importing economies. Clearly, as suggested by most recent studies on the macroeconomic effects oil shocks - including this one - modelling oil price shocks without taking into account the source of the shock is not a viable avenue.

A Appendix[16]

The idea of the coincident indicator model introduced by Stock and Watson (1991) is the notion that the comovements of many macroeconomic variables have a common element that can be captured by a single underlying, unobserved variable. The same idea is used here to try and gather relevant indicators for crude oil demand, and then combine them for a single underlying variable, which can then be used in the oil market VAR model.

Technically, I proceed by setting the indicator model in state-space form and estimating it with a Kalman filter.

The state-space form of the unobserved components (UC) model used here is:

$$Y_t = C_t' X_t + v_t \tag{2}$$

$$X_t = A_t X_{t-1} + F_t w_t \tag{3}$$

where equation (2) is the measurement equation, with Y_t an $n \times 1$ vector of observable variables (n is the number of variables, expressed in deviations from sample mean of month-on-month changes), C_t is an $(n \times (p + kn))$ matrix, X_t is a $(p + kn) \times 1$ vector of states, and v_t is an $n \times 1$ vector of error terms (with k lags to make it i.i.d). Equation (3) is the state equation, with a $(p + kn) \times (p + kn)$ matrix A_t, $(p + kn) \times (n + 1)$ matrix F_t and i.i.d error term w_t. In addition, the $(n + 1)x(n + 1)$ covariance matrix of the error terms of the state equation is also needed. In practice, in this application, the lag orders of both X_t (denoted by p) and v_t (denoted by k) are presumed to be two, and both F_t and A_t are time-invariant. Kalman filter smoothing is then used recursively to maximise a Gaussian log likelihood function to determine the parametre values for the model. The first element of the state vector X_t is then the final indicator.

[16] This is a brief description of state-space modelling with a Kalman filter. For more details, see, for example, Stock and Watson (1991).

B Appendix [17]

Let ϵ_t denote the structural VAR model innovations derived from equation (1). To construct structural impulse responses, one needs an estimate of the NxN matrix C in $u_t = C\epsilon_t$. Let $\Sigma_u = P\Lambda P$ and $C = P\Lambda^{1/2}$ such that C satisfies $\Sigma_u = CC'$. Then $C = BD$ (where B is a matrix of structural parametres obtained through a Choleski decomposition of the reduced form parameters) also satisfies $\Sigma_u = CC'$ for any orthonormal NxN matrix D.

It is possible to examine a wide range of possibilities for C by repeatedly drawing at random from the set **D** of orthonormal rotation matrices D. Following Rubio-Ramirez et. al (2005), I construct the set **C** of admissible models by drawing from the set **D** of rotation matrices and discarding candidate solutions for C that do not satisfy a set of a priori sign restrictions on the implied impulse response functions. The procedure follows these steps:

1. Draw an NxN matrix K of $NID(0,1)$ random variables. Derive the QR decomposition (to produce an orthonormal matrix and an upper-triangular matrix) of K such that $K = QR$ with the diagonal of R normalised to be positive.

2. Let $D = Q$. Compute impulse responses using the orthogonalisation $C = BD$. If all implied impulse response functions satisfy the sign restrictions, keep D. Otherwise, discard D.

3. Repeat the first two steps a large number of times, recording each D (and the corresponding impulse response functions) that satisfy the restrictions. The resulting **C** comprises the set of admissible structural VAR models.

[17] This section draws heavily on Rubio-Ramirez et al (2005) and Kilian and Murphy (2009).

C Appendix [18]

Define an impulse vector, which is a vector $a \in \mathbb{R}^N$ such that there exists some matrix A, where a is a column of A, such that $AA' = \Sigma_u$. Thus, the jth column of A represents the immediate impact, or impulse vector, of a one standard error innovation to the jth fundamental innovation, which is the jth element of ϵ_t. Furthermore, let C be the lower-triangular Choleski factor[19] of Σ_u and $Q = [q^{(1)}, ..., q^{(s)}]$ be an NxS matrix of orthonormal rows $q^{(i)}$, where S is the number of shocks to be identified in the model. Any impulse vector can then be written as, $a = Cq$, where q is the relevant column of Q, and $q = [q_1, ..., q_N]$, $\parallel q \parallel= 1$. Hence, q are the identifying weights to be determined. Following Uhlig (2005), the impulse responses for the impulse vector a can be written as a linear combination of the impulse responses to the Choleski decomposition of Σ_u as follows. Let $r_a(k)$ be the N dimensional impulse response at horizon k to the impulse vector a. The linear combination can then be written as:

$$r_a(k) = \sum_{i=1}^{N} q_i r_i(k) \tag{4}$$

where q_i is the i:th entry of q.

Next, define the penalty function f on the real line as $f(x) = 100x$ if $x > 0$ and $f(x) = x$ if $x \leq 0$. Let s_j be the standard error of variable j. Let $J_{S,+}$ be the index set of variables, for which identification of a given shock restricts the impulse response to be positive, and let $J_{S,-}$ be the index set of variables, for which identification restricts the impulse responses to be negative. To impose these sign restrictions, one solves for the weights q and thus $a = Cq$ by solving the minimisation problem:

$$q = \arg\min \Psi(Cq) \tag{5}$$

where the criterion function $\Psi(a)$ is given by:

$$\Psi(a) = \sum_{j \in J_{S,+}} \sum_{k=0}^{K} f(-\frac{r_{ja}(k)}{s_j}) + \sum_{j \in J_{S,-}} \sum_{k=0}^{K} f(\frac{r_{ja}(k)}{s_j})$$

The criterion function thus sums the penalties over the periods $k = 0, ..., K$ following the shock and over the indices of variables with positive ($J_{S,+}$) and negative ($J_{S.-}$) sign restrictions, respectively. The impulse responses are normalised by the standard error s_j of variable j. The penalty function is, of course, arbitrary.

Computationally, I implement this minimisation with a simplex algorithm, using the Rats econometric package.

[18] This section draws heavily on Mountford and Uhlig (2005). For further details, see Mountford and Uhlig (2005), Appendix A.

[19] The Choleski factorisation is not used for identification. It only serves as a computational tool, and any other factorisation would deliver the same results.

To identify two impulse vectors $[a^{(1)}, a^{(2)}]$, the first vector can be identified as detailed above, after which the vector pertaining to the second shock can be identified by replacing the minimisation problem in equation (5) with

$$q = \arg \min_{q' q^{(1)} = 0} \Psi(Cq) \tag{6}$$

i.e. by additionally imposing orthogonality to the first shock.

Likewise, if orthogonality to the first two shocks is required, the minimisation problem becomes

$$q = \arg \min_{q' q^{(1)} = 0, q' q^{(2)} = 0} \Psi(Cq) \tag{7}$$

In my model, I require the oil market shocks to be orthogonal to the US macroeconomic shocks, i.e., the oil market shocks are ordered causally first and the US macroeconomic shocks second. The computations are performed using a Bayesian approach as in Uhlig (2005). I take a number of draws from the posterior distribution. For each draw from the posterior of the VAR coefficients and the covariance matrix Σ_u, the shocks are identified as described above. Given the sample of draws for the impulse responses, confidence bands can be plotted.

D Appendix

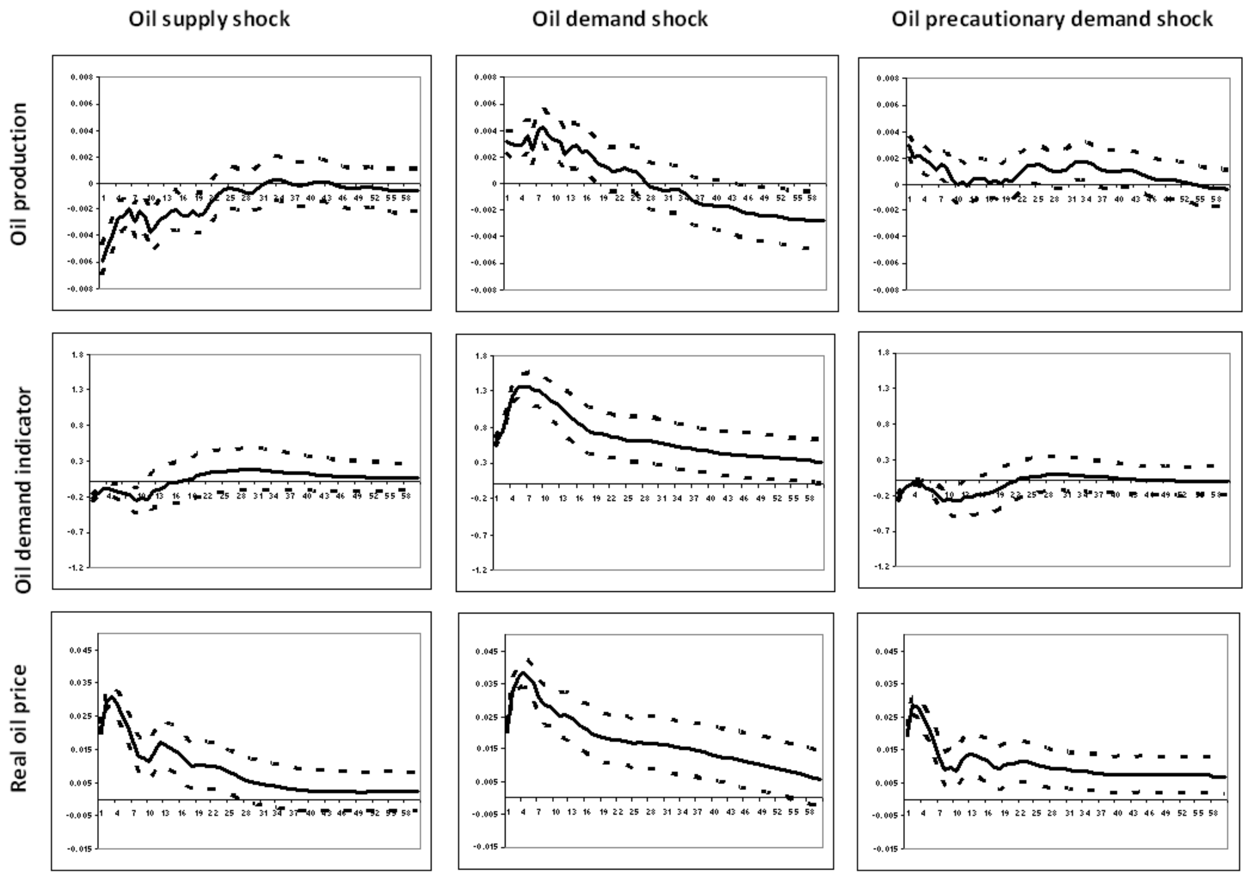

Figure 6: Median impulse responses in the benchmark model. Dashed lines indicate 16th and 84th quantiles based on 1000 accepted draws.

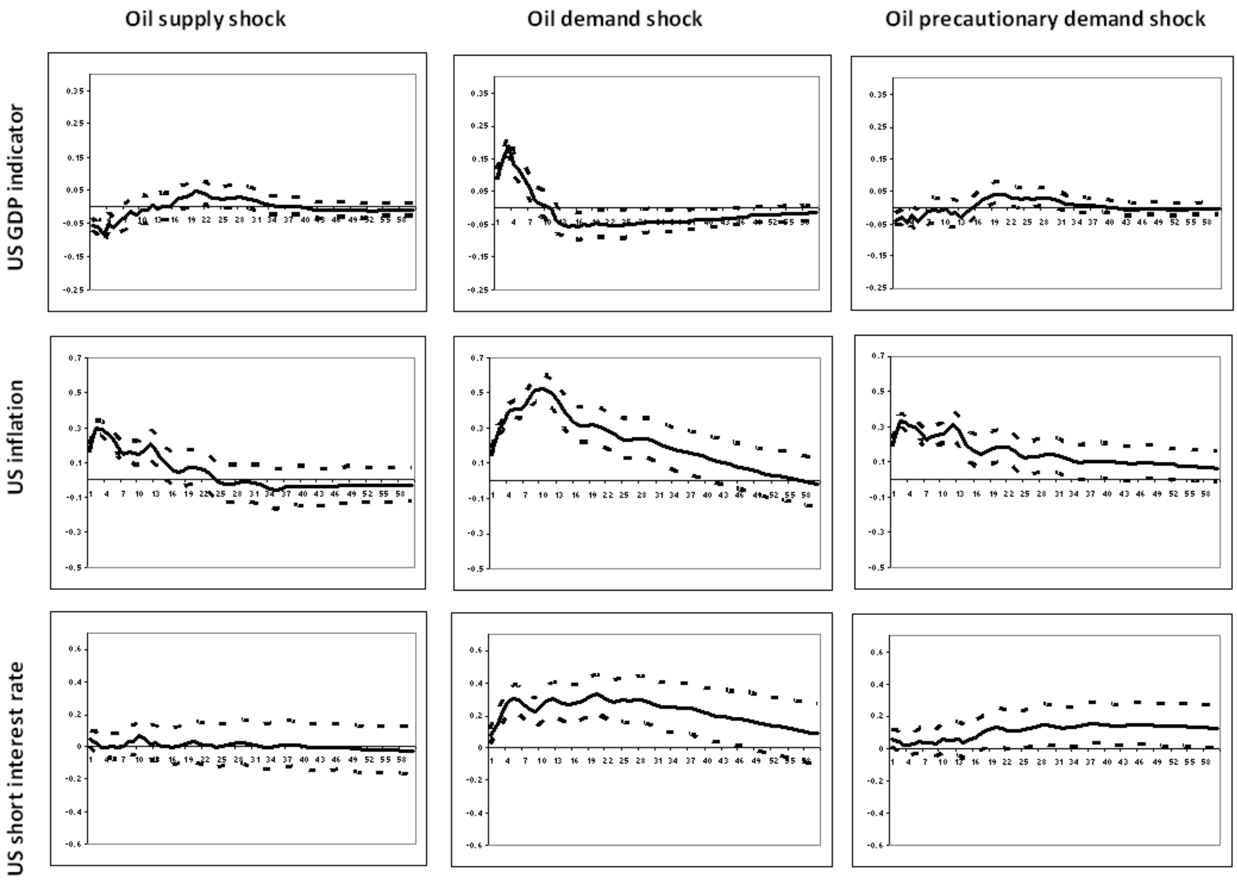

Figure 7: Median impulse responses in the benchmark model. Dashed lines indicate 16th and 84th quantiles based on 1000 accepted draws.

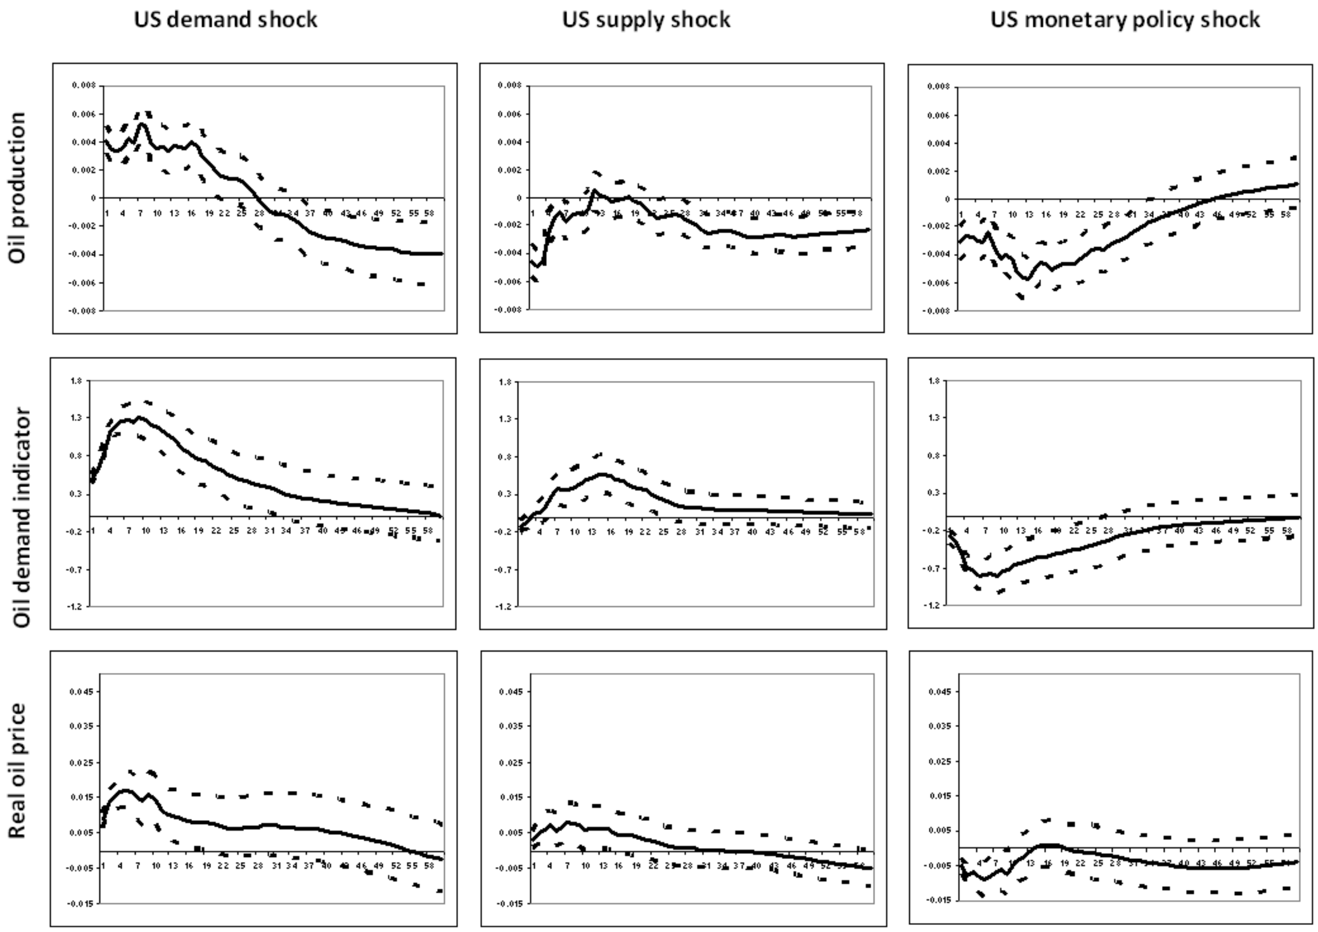

Figure 8: Median impulse responses in the benchmark model. Dashed lines indicate 16th and 84th quantiles based on 1000 accepted draws.

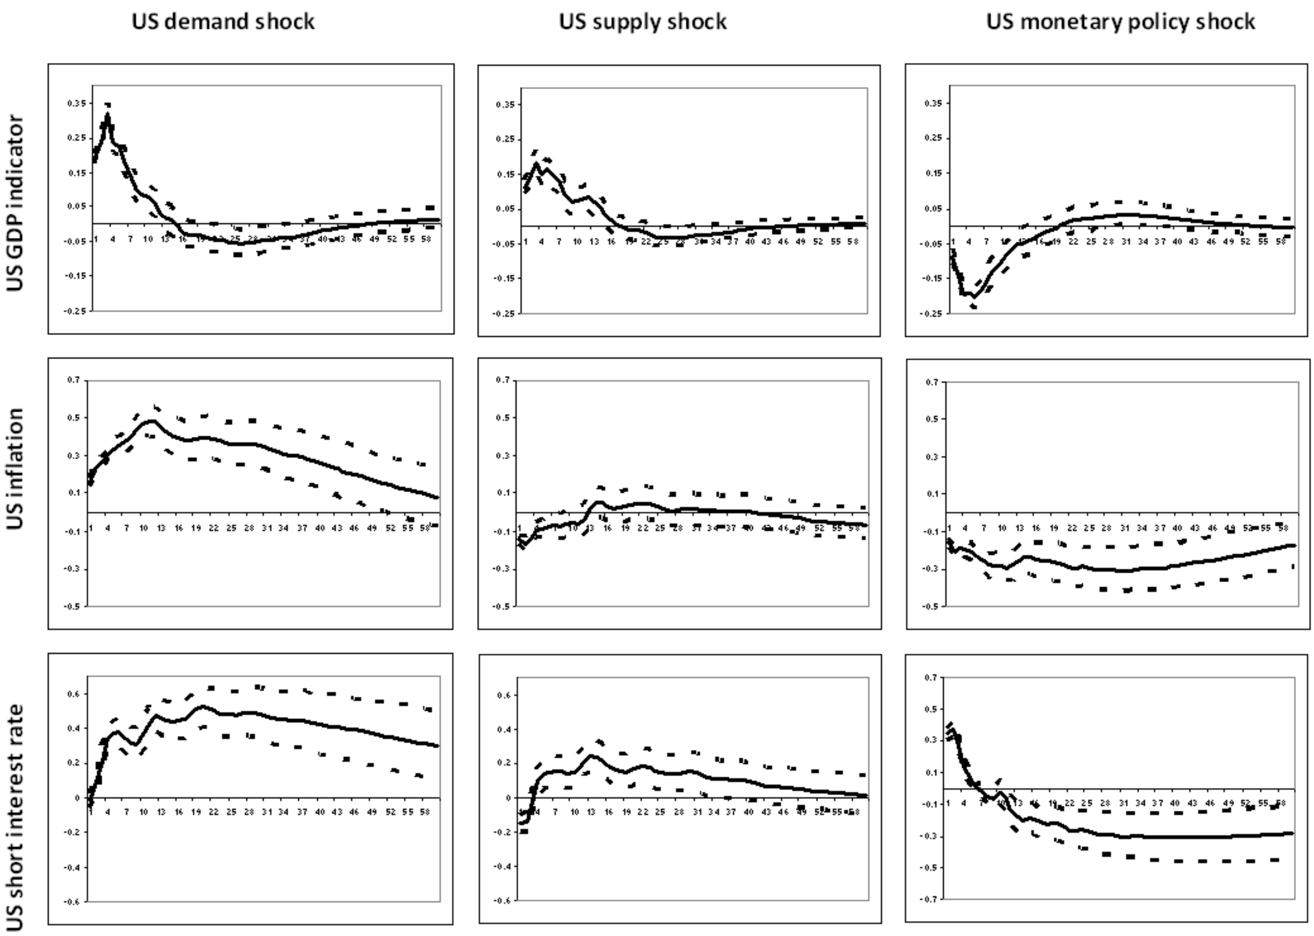

Figure 9: Median impulse responses in the benchmark model. Dashed lines indicate 16th and 84th quantiles based on 1000 accepted draws.

E Appendix

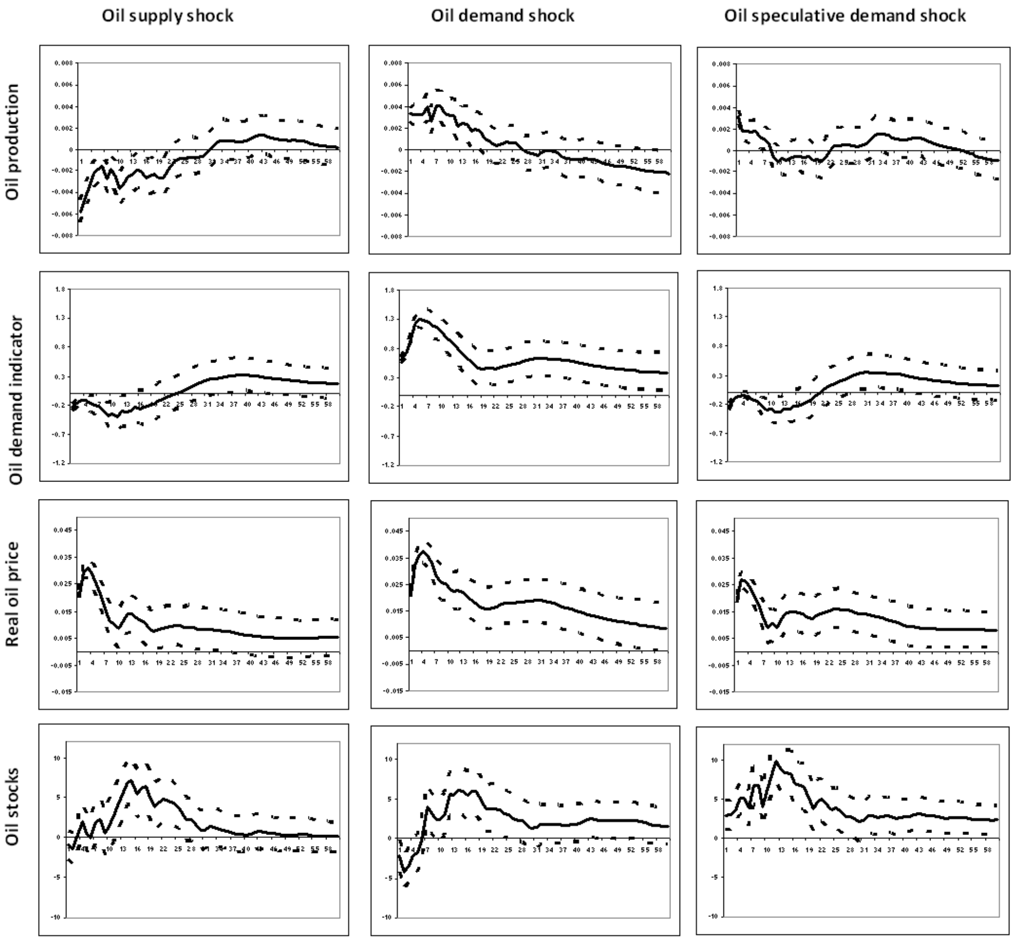

Figure 10: Selected median impulse responses in a 7-variable VAR with oil stocks. Dashed lines indicate 16th and 84th quantiles based on 1000 draws.

References

[1] Aastveit, K. A.: 2009, Modeling transmission of oil price shocks and monetary policy shocks in a data rich environment, mimeo

[2] Baumeister, C. and Peersman, G.: 2011, Time-varying effects of oil supply shocks on the US economy, mimeo.

[3] Baumeister, C. and Peersman, G.: 2009, Sources of the volatility puzzle in the crude oil market, mimeo, Ghent University

[4] Bernanke, B., Gertler, M. and Watson, M.: 1997, Systematic monetary policy and the effects of oil price shocks, Brookings Papers on Economic Activity 1, 91-157.

[5] Fry, R. and Pagan, A.: 2007, Some issues in using sign restrictions for identifying structural VARs. National Centre for Econometric Research Working Paper Number 14.

[6] Hamilton, J.: 2009, Causes and consequences of the oil shocks of 2007-08, Brooking Papers on Economic Activity 1, 215-259

[7] Hamilton, J. D. and Herrera, M.: 2004, Oil shocks and aggregate macroeconomic behavior: The role of monetary policy - a comment, Journal of Money, Credit and Banking 36 (2), 265-286.

[8] Kilian, L.: 2009, Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market, American Economic Review 99:3, 1053-1069.

[9] Kilian, L.: 2010, Oil price shocks, monetary policy and stagflation, in R. Fry, C. Jones and C. Kent (eds), Inflation in an Era of Relative Price Shocks, Sydney.

[10] Kilian, L. and Lewis, L. T.: 2010, Does the Fed respond to oil price shocks? The Economic Journal (forthcoming).

[11] Kilian, L. and Murphy, D.: 2009, Why agnostic sign restrictions are not enough: Understanding the dynamics of oil market VAR models. Journal of the European Economic Association (forthcoming).

[12] Kilian, L. and Murphy, D.: 2011, The role of inventories and speculative trading in the global market for crude oil, mimeo.

[13] Lippi, F. and Nobili, A.: 2009, Oil and the macroeconomy: a quantitative structural analysis, Temi di Discussione 704, Bank of Italy, Economic Research Department

[14] Lenza, M„ Pill, H. and Reichlin, L.: 2010, Monetary policy in exceptional times, European Central Bank Working Paper No. 1253, October 2010

[15] Lombardi, M. and Van Robays, I.: 2011, Do financial investors destabilize the oil price?, European Central Bank Working Paper No. 1346, June 2011

[16] Mountford, A. and Uhlig, H.: 2009, What are the effects of fiscal policy shocks? Journal of Applied Econometrics 24, 960-992.

[17] Peersman, G. and Van Robays, I.: 2009, Oil and the euro area economy, Economic Policy 24, 603-651.

[18] Rubio-Ramirez, J. F., Waggoner, D. F. and Zha, T.: 2010, Structural vector autoregressions: Theory of identification and algorithms for inference, The Review of Economic Studies 77, 665-696.

[19] Stock, J. H. and Watson, M. W.: 1991, A probability model of the coincident economic indicators, in K. Lahiri and G. Moore (eds), Leading Economic Indicators: New Approaches and Forecasting Records, Chapter 4.

[20] Uhlig, H.: 2005, What are the effects of monetary policy on output? Results from an agnostic identification procedure, Journal of Monetary Economics 52, 381-419.