A note on the identification and transmission of energy demand and supply shocks

Michelle, Gilmartin

9 January 2016

A note on the identification and transmission of energy demand and supply shocks

Michelle Gilmartin
University of Stirling

January 2016

Abstract

This paper proposes and implements a novel structural VAR approach for identifying oil demand and supply shocks. In this approach we search for two shocks in the context of a VAR model, which explain the majority of the k-step ahead prediction error variances of oil prices. Finally, we compare our approach with alternative identification schemes based on sign restrictions, and we show that the proposed methods is a useful tool for decomposing oil shocks.

Keywords: oil price shocks; demand and supply; Bayesian VAR; MCMC

JEL Classification: C11; E31; E32; Q43

1 Introduction

Identifying and measuring shocks in the market for energy is of paramount importance for policy-makers (e.g. a central bank trying to control expectations about consumer price inflation), producers and consumers in an economy, as well as investors in the stock and commodities markets. Therefore, it is no surprise that there exists a large academic literature employing structural vector autoregressions (SVARs). In this literature the estimated innovations from the VAR can be used to measure oil (and other energy) shocks to variables of interest such as output, price inflation, and a short term interest rate; see for example Bernanke, Gertler and Watson (1997). However, as it has been shown recently by authors such as Kilian (2009), Baumeister and Peersman (2013), and Lütkepohl and Netšunajev (2013), accurate measurement of oil shocks presupposes that we are able to decompose oil price dynamics into changes caused by demand or supply factors. This is because the intensity and transmission of oil demand shocks is different from that of oil supply shocks. Additionally, the effects of various oil shocks can have a diverse impact in an economy, so that, for example, monetary authorities need to identify the nature of the shock in order to react accordingly. A demand-driven increase in the price of oil will also result in an increase of the produced quantity of oil, which may be attributed to better prospects about global economic growth. In contrast, positive supply side shocks in the price of oil are expected to decrease the quantity of oil as well as output while at the same time causing a stagflation effect by pushing total (headline) consumer inflation to increase.

In this paper we reassess the issue of measuring empirically demand and supply oil shocks using vector autoregressions. Given the evident nonlinearities between oil prices and many macroeconomic variables such as GDP (see for instance Ravazzolo and Rothman, 2012, and references therein), we follow Clark and Terry (2010) and Baumeister and Persman (2013) and our preferred specification is a VAR on energy prices, quantity of energy, GDP and CPI, which features time-varying coefficients and stochastic volatility. We examine the implications of an identification scheme, first proposed by Uhlig's (2003) and materialized by Barsky and Sims (2011), which decomposes energy price shocks into two components, orthogonal to each other, such that both energy price shocks maximize the forecast error variance of energy prices. We find a clear interpretation of the first of these two shocks as a supply-side energy shock, resulting in a change of energy prices and CPI on the one direction and a change of the quantity of energy and GDP in the opposite direction. The second shock can be interpreted as demand-side shock which results in a change in all four variables in the system toward the same direction. We find that this identification scheme is quite similar to using sign restrictions (Baumeister and Persman, 2013), however, it gives much more reasonable estimates of the impulse responses which are in accordance with the reccomndations in Kilian and Murphy (2012).

So far there are several existing methods for implementing a decomposition of VAR innovations into demand-driven and supply-driven shocks. The vast majority of these methods have been applied to the market for crude oil, however, they apply for general energy markets. Kilian (2009) has proposed to use exclusion restrictions based on a recursively identified VAR for monthly measurements of oil supply, global demand, and oil demand. He assumes that the short-run supply curve of oil is vertical, by sorting oil supply first in the VAR and assuming that this variable is exogenous and does not respond to innovations in global demand or oil-specific demand. Baumeister and Peersman (2013), motivated by the fact that oil-supply cannot remain exongenous over the quarter, when using quarterly data, they propose to use sign restrictions. Kilian and Murphy (2012) go one step further and show that sign restrictions alone result in overestimation of the response of variables to supply shocks and underestimation of responses to demand shocks, thus suggesting to impose additional restrictions on the oil supply elasticities on impact.

All these methods impose a great deal of restrictions in the empirical econometric models. Even though such restrictions may be driven from economic theory, it doesn't necessarily mean that they are or that they should be supported by evidence in the data (the likelihood). Such assumptions are very beneficial when estimating time-invariant VARs. However, given the evident changing dynamics of the US economy (e.g. Great Inflation, Great Moderation, Great Recession) as well as the volatile character of oil price shocks, it might be dangerous to impose such structure when identifying oil shocks. For instance, both demand and supply elasticities might have changed dramatically over the course of the past 40 years, so it is less clear how to adapt the important contributions of Kilian and Murphy (2012) in the time-varying parameters framework we are working with, in order to impose bounds on the elasticities at each point in time[1].

In this paper we are not arguing that there are no caveats when applying Uhlig's (2003) "data-based" identification scheme in an energy market VAR. For example, the identification scheme we use makes the assumption that only two shocks affect oil prices at all forecast horizons. This assumption could potentially be dismissed by some economists, despite the fact that we show that these two shocks have an interpretation as demand and supply shocks. Nevertheless, we show that identification of structural shocks is achieved with a minimal set of restrictions compared to other methods, thus allowing an alternative interpretation of energy price shocks. When we compare our results to demand and supply shocks identified using sign restrictions on the impact matrix, we find a surprisngly high degree of similarity of the shape of impulse responses, but the magnitutes can be quite different. However, there are some differences especially in the measurement of the initial impact of an oil price increase which is due to a supply shock, and about the measurement of the pass-through of oil prices to CPI. In this paper we assess all this evidence by comparing impulse responses and we discuss the implications for policy-making.

The next Section introduces the estimated reduced-form time-varying parameter VAR model we estimate, and identification of structural shocks. In Section 3 we present all the empirical evidence including time-varying impulse resposnes as well as time-varying forecast error variance decompositions attributed to a demand and a supply shock. Section 4 concludes the paper.

2 Empirical Methodology

2.1 Time-varying coefficients VAR with stochastic volatility

Following Clark and Terry (2010), our starting point is a p-lag vector autoregression with time-varying parameters and stochastic volatility estimated for global oil production (q_t^{oil}), the real

[1] Such issues are explained in detail in the seminal contribution of Baumeister and Persman (2013) who are the first to examine oil market VARs with structural instabilities.

(deflated using US CPI) US refiners' acquisition cost of imported crude oil (p_t^{oil}), real GDP (gdp_t), and the consumer price index (cpi_t)[2]. The model takes the following form

$$y_t = X_t' \beta_t + A_t^{-1} \Sigma_t \varepsilon_t,$$

where $y_t = \left[q_t^{oil\prime}, p_t^{oil\prime}, gdp_t', cpi_t',\right]'$, $X_t = I_4 \otimes [1, y_{t-1}, ..., y_{t-p}]$, and $\varepsilon_t \sim N(0, I_4)$. The vector β_t contains all the VAR coefficients for the intercept and p-lags, where we set $p = 4$ to allow our VAR to capture sufficient dynamics (see Hamilton and Herrera, 2004). A_t is a lower triangular matrix with ones on the diagonal with element $\alpha_{ij,t}$, $i = 2, ..., 4$, $j = 1, ..., i-1$, and Σ_t is a diagonal matrix consisting of the standard deviations σ_{1t}, such that $\Omega_t = A_t^{-1} \Sigma_t \left(\Sigma_t A_t^{-1}\right)'$ is the full VAR covariance matrix at time t.

Using the standard practice in this literature (Clark and Terry, 2010, and references therein) the coefficients β_t, $\alpha_t = (\alpha_{21,t}, ..., \alpha_{ij,t})$, and $\sigma_t = (\sigma_{1t}, ..., \sigma_{4t})$ follow random walks of the form

$$\begin{aligned} \beta_t &= \beta_{t-1} + \eta_t^{\beta}, \\ \alpha_t &= \alpha_{t-1} + \eta_t^{\alpha}, \\ \log \sigma_t &= \log \sigma_{t-1} + \eta_t^{\sigma}, \end{aligned}$$

where $\eta_t^{\beta} \sim N(0, Q)$, $\eta_t^{\alpha} \sim N(0, S)$, and $\eta_t^{\sigma} \sim N(0, W)$ are state errors uncorrelated with each other, as well as ε_t, at all leads and lags. The covariance matrix Q is of dimension $k \times k$ where $k = (p+1) \times 4 \times 4$, S is a 4×4 matrix, and W is a 4×4 diagonal matrix.

It is straightforward to show that such a structure implies that the equations above are conditionally Gaussian linear models, that is, one can estimate any of the time-varying parameters $(\beta_t, \alpha_t, \log \sigma_t)$ conditional on all other parameters using the Kalman filter/smoother. Such conditioning is natural in Bayesian Markov Chain Monte Carlo (MCMC) methods, and in particular the Gibbs sampler which is the preferred method in this paper. The reader is referred to Koop and Korobilis (2010) for exact estimation details. Here it suffices to note that since we use Bayesian methods prior distributions over all time-varying and time-invariant model parameters need to be carefully specified. In this paper we follow Baumeister and Peersman (2013) and others and estimate a VAR in the pre-sample period 1947Q1-1973Q4 using OLS, and we use these parameter estimates as starting values (prior hyperparameters) for the sample of interest 1974Q1-2012Q4.

2.2 Identifying supply and demand shocks

Once the TVP-VAR parameters are estimated, we can obtain the structural, time-varying parameter, Vector Moving Average form of the VAR which is

$$y_t = C_t(L) H_{0,t} \varepsilon_t,$$

where $C_t(L) = \left(I - B_{1t}L - B_{2t}L^2 - ... - B_{pt}L^p\right)^{-1}$, B_{it}, for $i = 1, ..., p$, are 4×4 VAR coefficient matrices obtained by rearranging the elements of the vector β_t, and $H_{0,t}$ is the impact matrix which satisfies $H_{0,t} H_{0,t}' = \Omega_t$. There are many matrices which satisfy this condition, for

[2]All variables are transformed to growth rates by taking first log-differences and multiplying by 100, with the exception of GDP which is annualized quarter-on-quarter growth rates (i.e. multiplied by 400).

example one solution is to set $H_{0,t} = chol\left(\Omega_t\right) = \Sigma_t A_t^{-1}$ which gives the standard recursive identification that has been used extensively in the literature. Such $H_{0,t}$ is lower triangular, this implies that exclusion restrictions where the preceding variables in y_t are fully exogenous to the succeeding variables, so that the first variable y_t is not affecting any other variable on impact of a structural shock, while the last variable affects all preceeding variables. Kilian (2009) has used such exclusion restrictions in order to identify oil demand and supply shocks in the global market[3]. When using quarterly data it is less clear why a specific variable in the VAR would not respond contemporaneously to a shock in any other variable, thus making this sort of identification scheme less attractive.

Alternative impact matrices can be obtained by rotating any solution $H_{0,t}$ using an arbitrary orthonormal matrix D ($D\prime D = I$). In this case it holds that $\widetilde{H}_{0,t} = H_{0,t}D'$ is also a valid impact matrix since $\widetilde{H}_{0,t}\widetilde{H}_{0,t}' = H_{0,t}D'DH_{0,t}' = H_{0,t}H_{0,t}' = \Omega_t$. Note that in practice we go the other way around, so that we have an estimate of Ω_t from the reduced form VAR, and there are infinite matrices $\widetilde{H}_{0,t}$ (or D) that one can identify. In this paper we present two such cases. The main identification scheme first proposed by Barsky and Sims (2011) and Uhlig (2003) which we will adapt to identify oil shocks over time. We also explain briefly the procedure for imposing sign restrictions on the impact matrix, similar (but not identical) to the one used in Baumeister and Persman (2013), which will serve as a benchmark for evaluation of our results.

When imposing sign restrictions, D is identified as a rotation matrix which covnerts $H_{0,t}$ into a matrix $\widetilde{H}_{0,t}$ that satisfies all sign restriction conditions, e.g. a positive oil supply shock causes output to contract, thus causing a negative sign on the respective element of $\widetilde{H}_{0,t}$.

3 Empirics

3.1 Data and model estimates

We use a quartery dataset which is an updated version of the data used in Baumeister and Peersman (2013), available at http://www.aeaweb.org/aej/mac/data/2011-0065_data.zip. The variables are the world oil production, real refiner's acquisition cost, real GDP and CPI, and are observed for the period 1947Q1-2012Q4. The first part of the sample, 1947Q1-1973Q4, is used to "train the data" by estimating a time-invariant VAR using least squares, and then use the OLS estimates as starting values for the time-varying parameter VAR estimated over the period of interest, 1974Q1-2012Q4. The TVP-VAR (and, of course the VAR in the training sample) has four lags. Our results are quite robust to the choice of lags, meaning that the general shapes of the impulse responses do not change dramatically.

[3] Unlike this paper, or Baumeister and Peersman (2013), Kilian (2009) also uses an index of global economic activity which plays a crucial role in the decomposition of oil demand and supply shocks.

3.2 Impulse responses: supply vs demand shocks

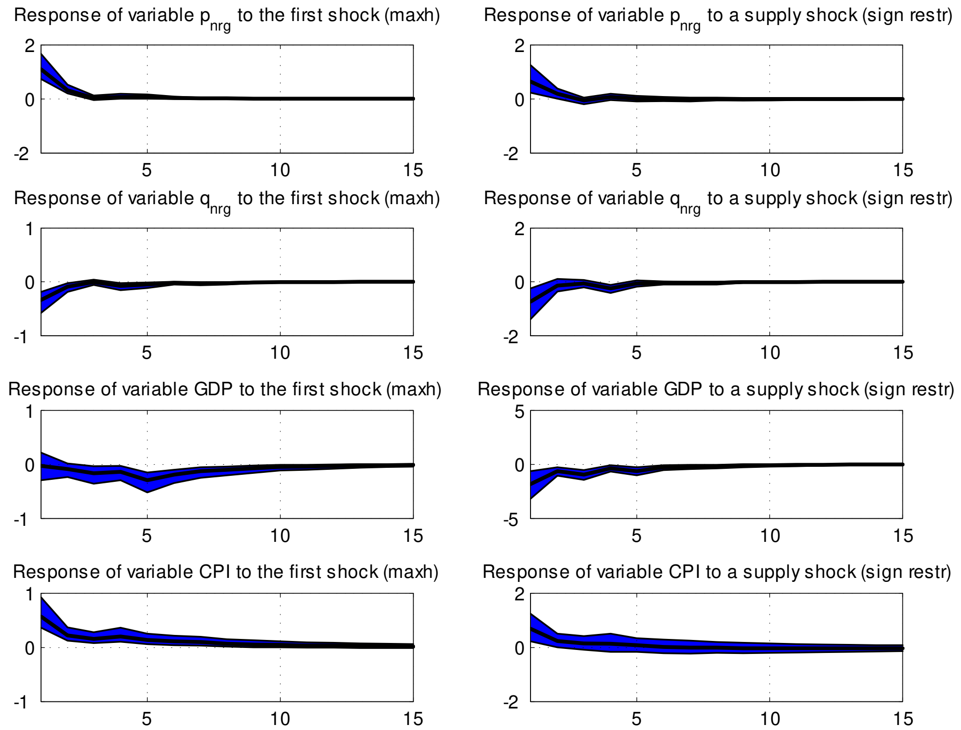

Median impulse responses (solid line) with 68% band (shaded area) of all four variables in the VAR, in 1977Q2 - supply shocks.

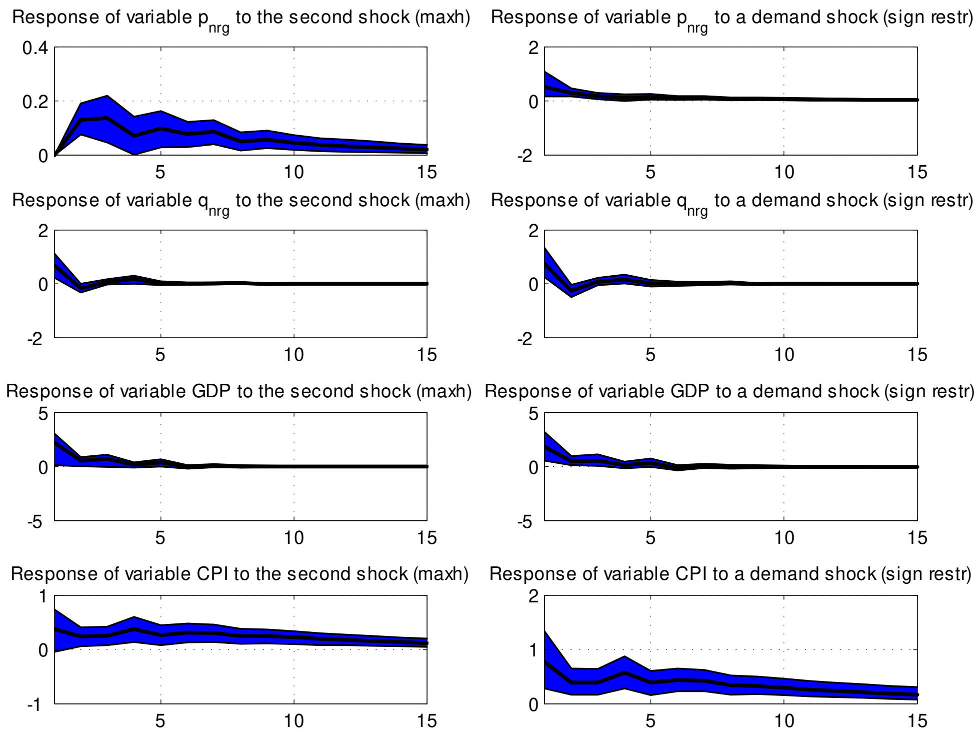

Median impulse responses (solid line) with 68% band (shaded area) of all four variables in the VAR, in 1977Q2 - demand shocks.

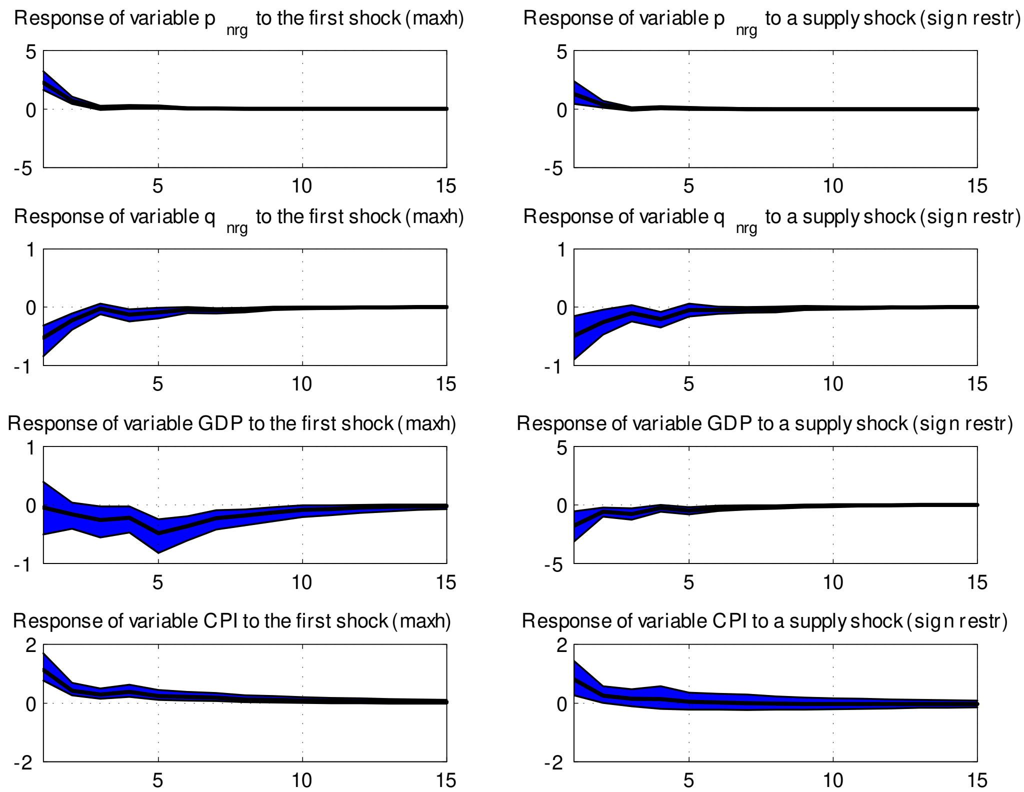

Median impulse responses (solid line) with 68% band (shaded area) of all four variables in the VAR, in 1982Q2 - supply shocks.

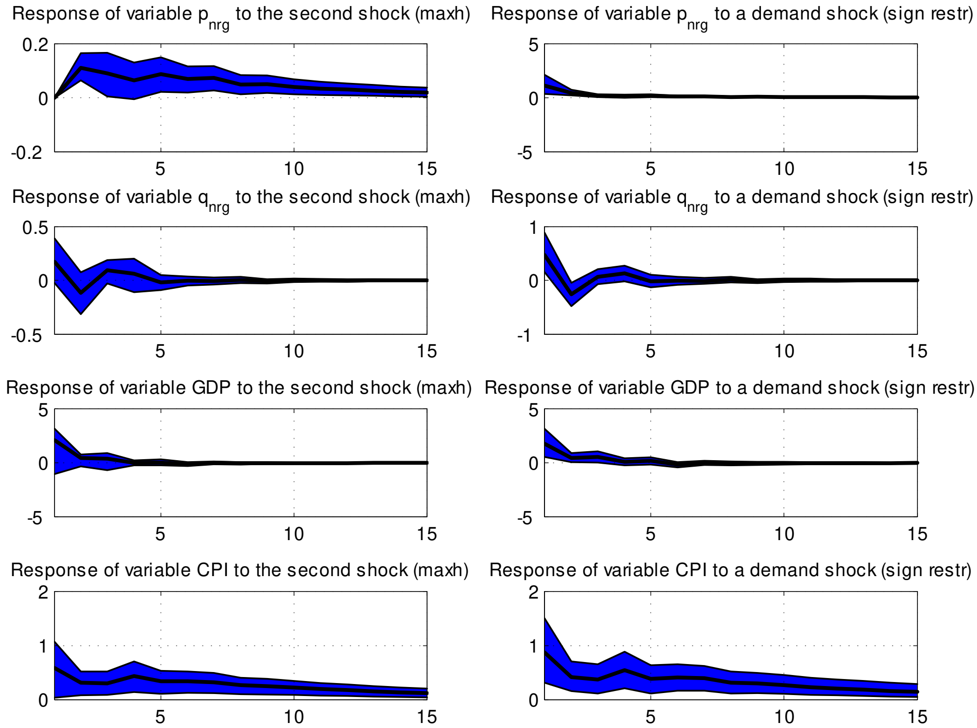

Median impulse responses (solid line) with 68% band (shaded area) of all four variables in the VAR, in 1982Q2 - demand shocks.

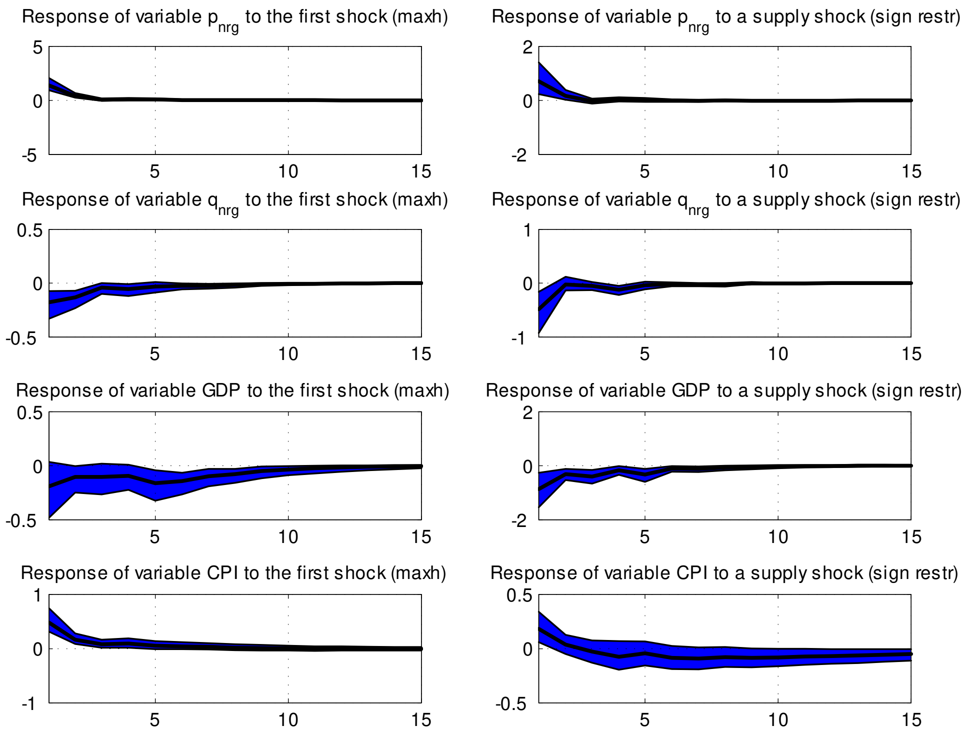

Median impulse responses (solid line) with 68% band (shaded area) of all four variables in the VAR, in 1994Q4 - supply shocks.

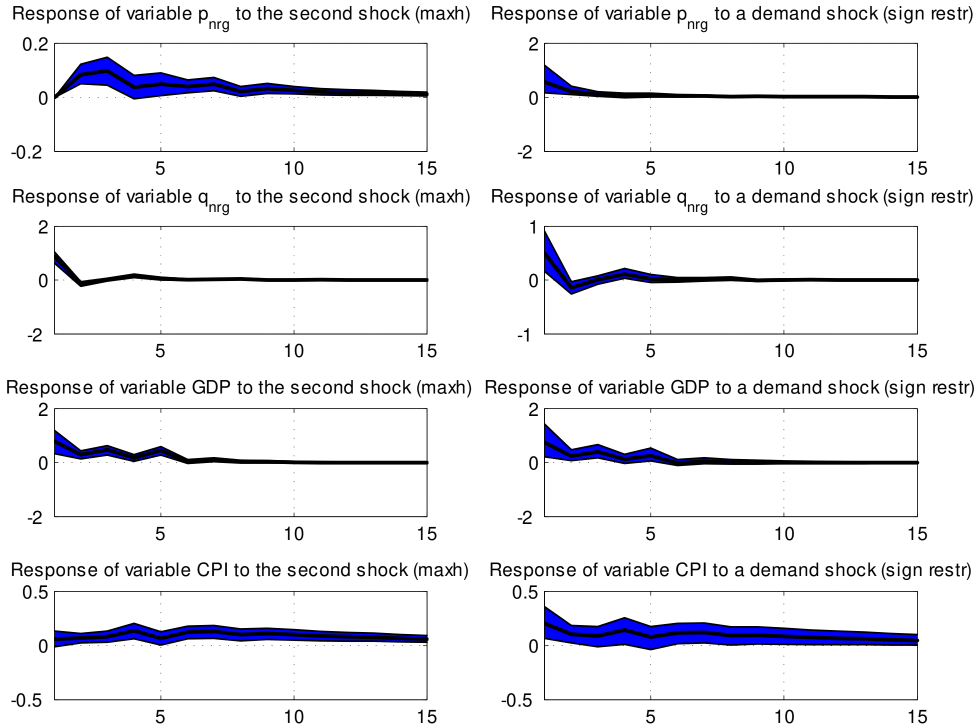

Median impulse responses (solid line) with 68% band (shaded area) of all four variables in the VAR, in 1994Q4 - demand shocks.

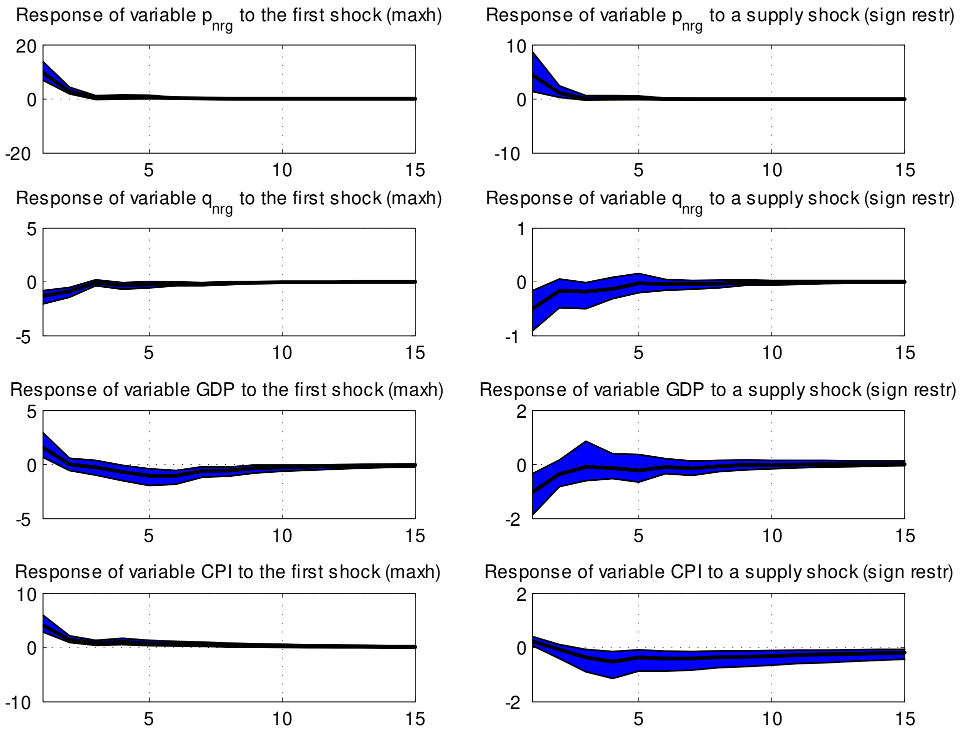

Median impulse responses (solid line) with 68% band (shaded area) of all four variables in the VAR, in 2009Q1 - supply shocks.

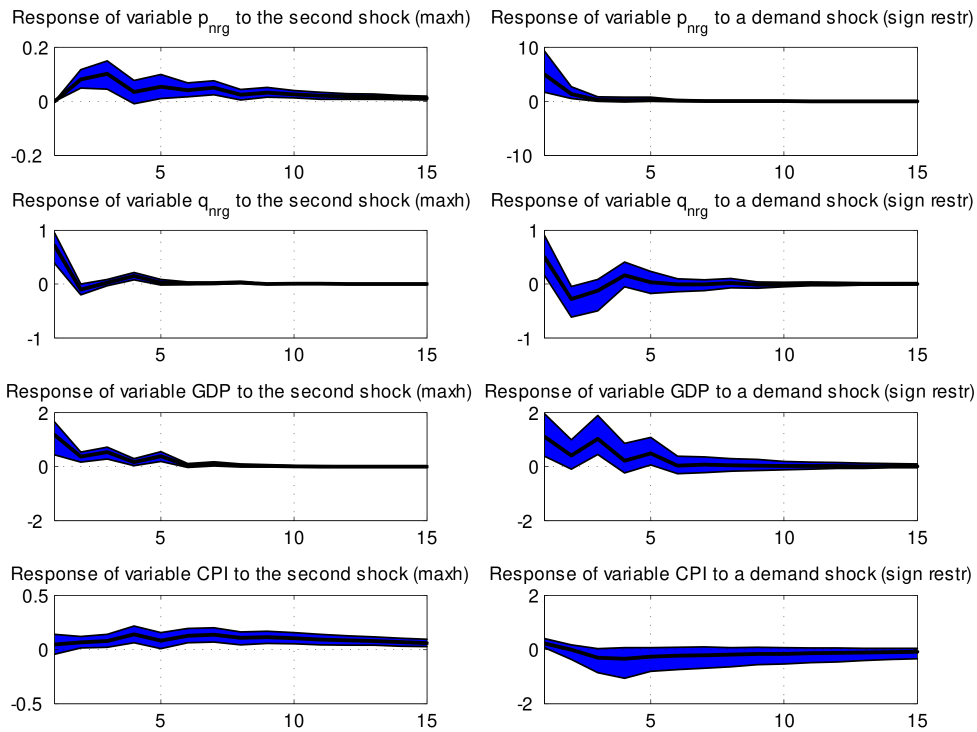

Median impulse responses (solid line) with 68% band (shaded area) of all four variables in the VAR, in 2009Q1 - demand shocks.

3.3 Forecast error variance decompositions

Another positive aspect of identifying demand and supply shocks in a time-varying parameter VAR is that we are able to obtain separate forecast error variance decompositions (FEVDs) for the oil demand and for the oil supply shock, over time. That is, one can estimate the fraction of the variance of the forecast error of, say, GDP which is attributed in a demand-side induced in the price of oil (as opposed to supply-driven increases) for any specific forecast horizon of interst.

This evidence is provide in Figure XXX. The left panels show shows the posterior median of the time-varying FEVDs of demand-side shocks to oil price, oil quantity, GDP and CPI, while the right panel of this figure shows the supply-driven FEVDs for the same set of variables in our TVP-VAR.

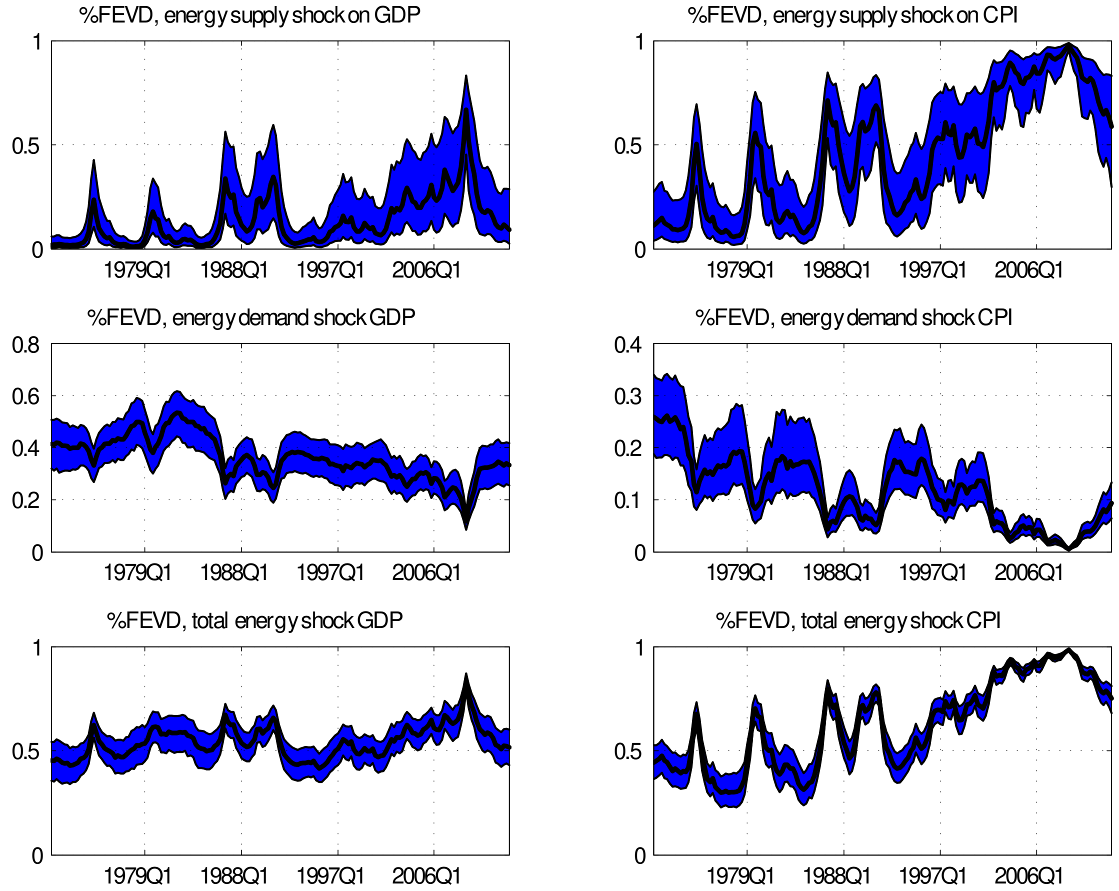

Medians (solid line) and 16th and 84th bands (shaded area) of the posterior of the time-varying forecast error variance at the 20 quarter ahead horizon, of GDP (left panel) and CPI (right panel) which is due to a supply, a demand, and a supply+demand change in oil prices, respectively. Numbers in the vertical axis denote the percentage (%) of forecast error variance explained.

4 Conclusions

In this note we propose an alternative method for decomposing demand and supply shocks based on a statistical identification method that searches for two (or more) shocks that maximize the forecast error variance at a given horizon.

5 References

Baumeister, Christiane, and Gert Peersman (2013). Time-varying effects of oil supply shocks on the US economy. *American Economic Journal: Macroeconomics*, 5(4), 1-28.

Barsky, Robert B., and Eric R. Sims (2011). News shocks and business cycles. *Journal of Monetary Economics*, 58, 273-289.

Bernanke, Ben, Mark Gertler, and Mark W. Watson (1997). Systematic monetary policy and the effects of oil price shocks. *Brookings Papers on Economic Activity*, 1, 91-142.

Clark, Todd E., and Stephen J. Terry (2010). Time variation in the inflation passthrough of energy prices. *Journal of Money, Credit and Banking*, 42(7), 1419-1433.

Hamilton, James D., and Ana María Herrera. (2004). Comment: Oil shocks and aggregate macroeconomic behavior: The role of monetary policy. *Journal of Money, Credit, and Banking*, 36(2), 265-86.

Kilian, Lutz (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.

Kilian, Lutz and Daniel P. Murphy (2012). 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.

Koop, Gary and Dimitris Korobilis (2011). UK macroeconomic forecasting with many predictors: Which models forecast best and when do they do so? *Economic Modelling*, 28, 2307-2318.

Koop, Gary and Dimitris Korobilis (2013). Large time-varying parameter VARs. *Journal of Econometrics*, 177, 185-198.

Koop, Gary and Dimitris Korobilis (2014). A new index of financial conditions. *European Economic Review*, 71, 101-116.

Koop, Gary, Leon-Gonzales, Roberto, and Rodney Strachan (2009). On the evolution of the monetary policy transmission mechanism. *Journal of Economic Dynamics and Control*, 33, 997-1017.

Koop, Gary, Pesaran, M. Hashem and Simon M. Potter (1996). Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics,* 74, 77-118.

Korobilis, Dimitris (2013a). Assessing the transmission of monetary policy shocks using time-varying parameter dynamic factor models. *Oxford Bulletin of Economics and Statistics*, 75, 157–179.

Korobilis, Dimitris (2013b). VAR forecasting using Bayesian variable selection. *Journal of Applied Econometrics*, 28, 204-230.

Lütkepohl, Helmut and Alexei Netšunajev (2013). Disentangling demand and supply shocks in the crude oil market: how to check sign restrictions in structural VARs. *Journal of Applied Econometrics*. DOI: 10.1002/jae.2330.

Lütkepohl, Helmut (2006). *New introduction to multiple time series analysis*. Springer.

Ravazzolo, Francesco and Phillip Rothman (2012). Oil and US GDP: A real-time out-of-sample examination. *Journal of Money, Credit and Banking*, 45(2-3), 449-463.

Uhlig, Harald (2003).What drives GNP? Unpublished manuscript, Euro Area Business Cycle Network.

A Appendix

In this Appendix we describe the estimation algorithm and the priors used to obtain the results in this paper. We also provide additional results about model estimation and identification of oil demand and supply shocks which, to our opinion, can further assist the reader to understand the mechanics behind our methodology.

A.1 Estimation of the model

Consider the time-varying parameter VAR

$$y_t = X_t' \beta_t + A_t^{-1} \Sigma_t \varepsilon_t,$$

with the additional equations describing the evolution of the parameters β_t, $\alpha_t = (\alpha_{21,t}, ..., \alpha_{ij,t})$, $i = 2, ..., 4$, $j = 1, ..., i-1$, and $\sigma_t = (\sigma_{1t}, ..., \sigma_{4t})$ using random walks of the form

$$\begin{aligned} \beta_t &= \beta_{t-1} + \eta_t^{\beta}, \\ \alpha_t &= \alpha_{t-1} + \eta_t^{\alpha}, \\ \log \sigma_t &= \log \sigma_{t-1} + \eta_t^{\sigma}, \end{aligned}$$

where $\eta_t^{\beta} \sim N(0, Q)$, $\eta_t^{\alpha} \sim N(0, S)$, and $\eta_t^{\sigma} \sim N(0, W)$ are state errors uncorrelated with each other, as well as ε_t, at all leads and lags. The covariance matrix Q is of dimension $k \times k$ where $k = 4 \times (4 \times 4 + 1) = 68$ in our application, S is a 4×4 matrix, and W is a 4×4 diagonal matrix.

We define the following initial conditions on the time-varying parameters

$$\begin{aligned} \beta_0 &\sim N\left(E\left(\beta^{OLS}\right), 4 \times var\left(\beta^{OLS}\right)\right), \\ \alpha_0 &\sim N\left(E\left(\alpha^{OLS}\right), 4 \times var\left(\alpha^{OLS}\right)\right), \\ \log \sigma_0 &\sim N\left(\log\left(E\left(\alpha^{OLS}\right)\right), 4 \times I_5\right). \end{aligned}$$

where $E\left(x^{OLS}\right)$ denotes the OLS estimate of a parameter x, and $var\left(x^{OLS}\right)$ its covariance

Finally, the state covariances Q, S and W have priors

$$\begin{aligned} Q &\sim iW\left(64 + 1, 0.0001 \times var\left(\beta^{OLS}\right)\right), \\ S_j &\sim iW\left(j + 1, 0.01 \times var\left(\alpha^{OLS}\right)_j\right), j = 1, 2, 3 \\ W_{ii} &\sim iG(8, 0.001). \end{aligned}$$

Note that for estimation purposes S needs to have a specific block diagonal structure, with the diagonal consisting of the 1×1, 2×2 and 3×3 submatrices S_1, S_2 and S_3, respectively. Estimation is implemented using Markov Chain Monte Carlo (MCMC) and in particular the Gibbs sampler; the reader is referred to Koop and Korobilis (2010) for details.

A.2 Assessment

Figure XXX plots the time-varying standard deviations, i.e. the square root of the diagonal elements of Ω_t. These estimates over time are the ones expected from previous experience with similalr VARs. On average (i.e. if we take the mean over all time periods t of the posterior median of Ω_t) the time-varying parameter VAR generates smaller standard deviations than its constant counterpart estimated with OLS. This evidence serves as a good rule-of-thumb approach to assess the fit of time-varying parameter VARs, given that calculation of marginal likelihoods and model selection is computationally cumbersome - if not infeasible; see the discussion in Korobilis (2013).

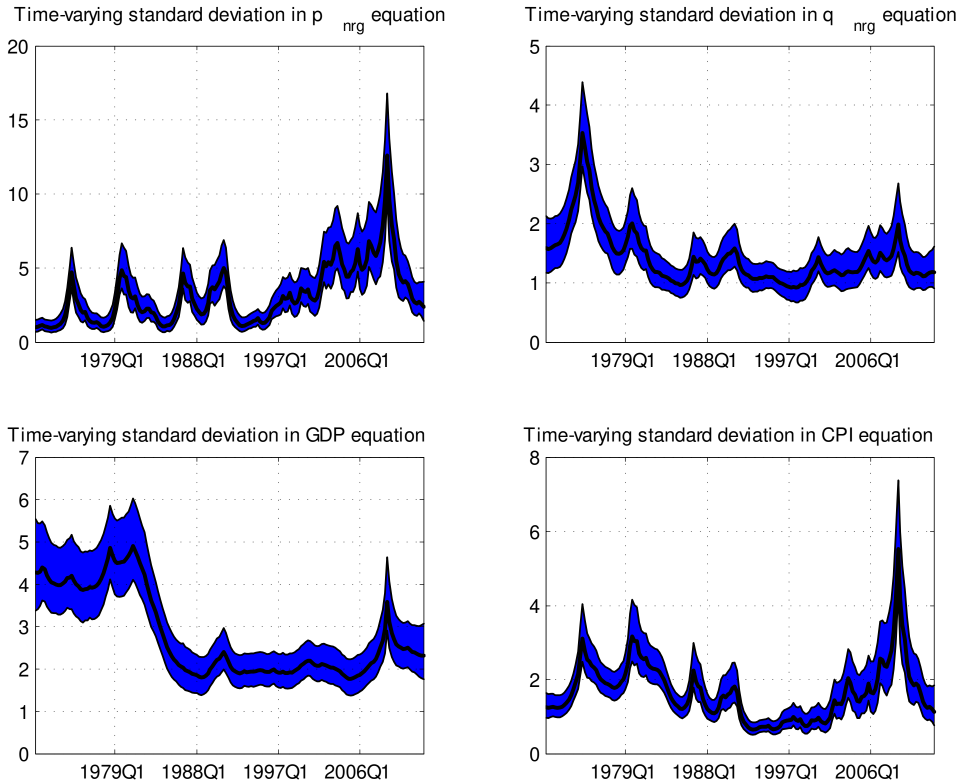

Medians (solid line) and 16th and 84th bands (shaded area) of the posterior of the time-varying standard deviations from the TVP-VAR.

Figure XXX plots the time-varying correlations between the price of oil and GDP (upper panel), and the price of oil and CPI (lower panel). Again, these results are the ones expected from experience and common sense. Firsth, there is a clear negative relationship between oil prices and GDP which collapses during the 2007-2009 financial crisis (where, eventually, both oil prices and GDP collapsed in 2008). Second, there is more of a negative relationship between oil prices and CPI with regular "spikes" where correlation becomes positive (e.g. around the recent financial crisis). Oil price inflation (nominal) on has been increasing at a very aggressive rate during the post-Great Moderation period, which is a period where fluctuations on CPI inflation have been quite moderate, thus explaining the negative (on average) correlation between real oil prices and CPI captured by the TVP-VAR.

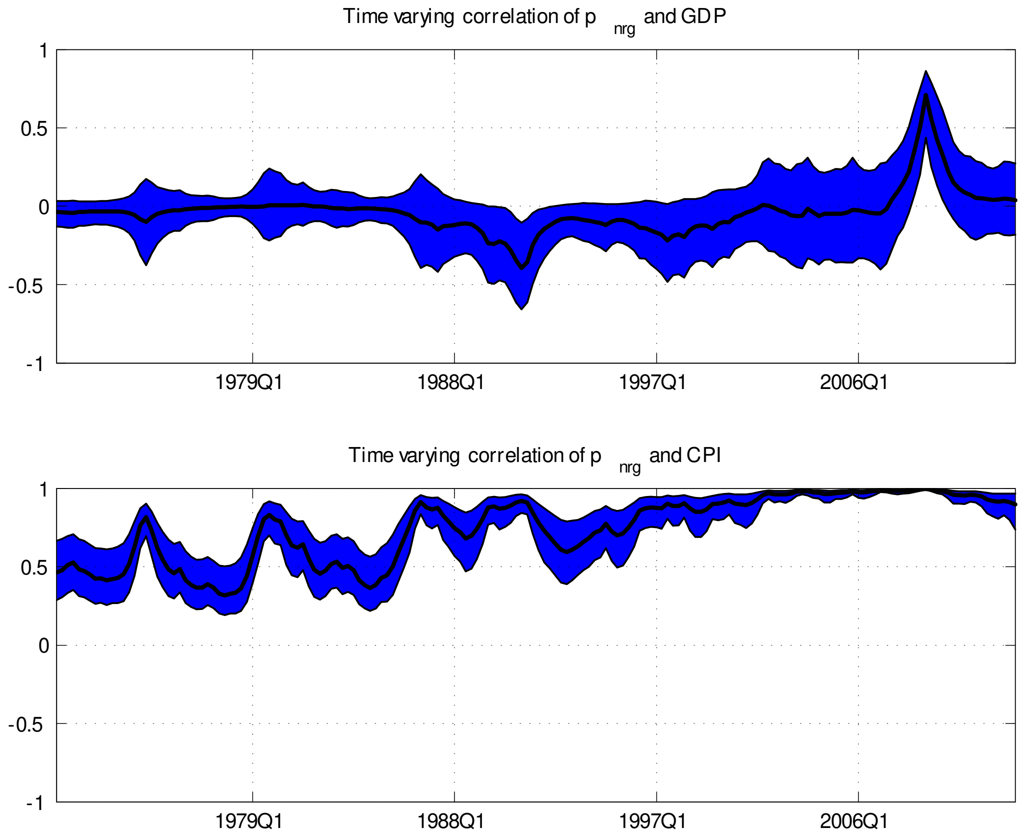

Medians (solid line) and 16th and 84th bands (shaded area) of the posterior of time-varying correlations between oil price and gdp (upper panel), and oil price and cpi (lower panel).

A.3 Further results

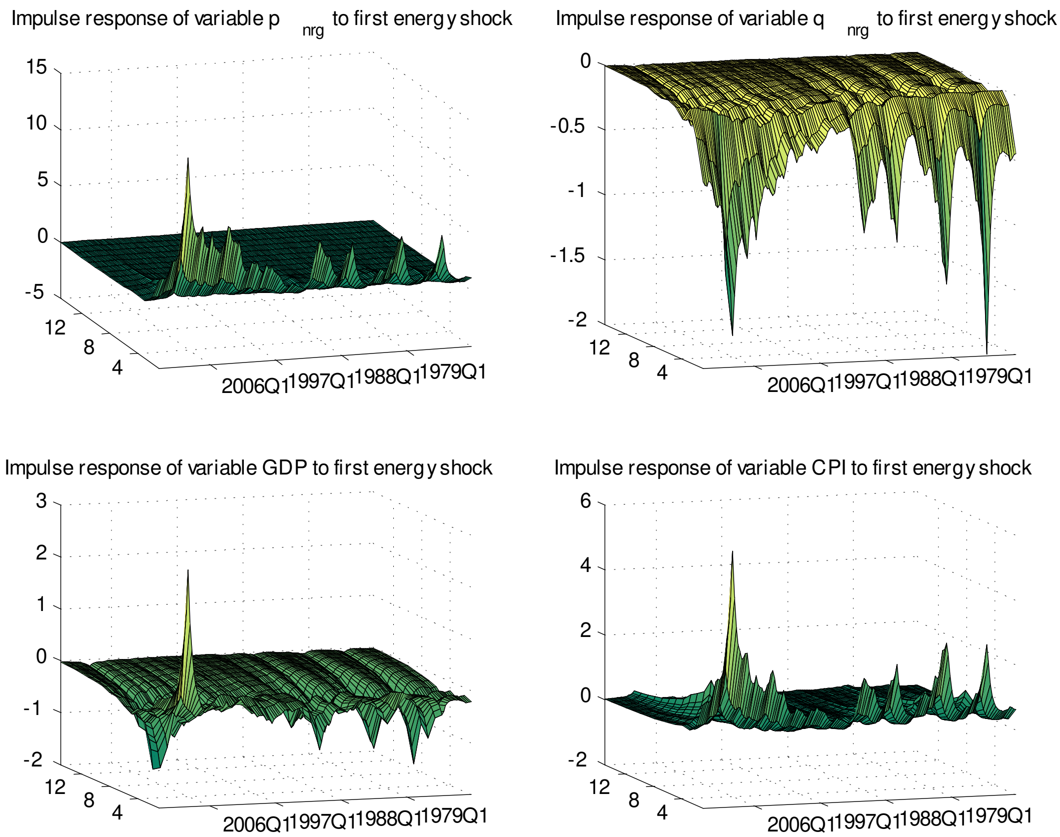

Posterior medians of impulse responses for all 20 horizons and all time-periods (1970Q1-2012Q4). Shock is the first identified ("supply") shock from the decomposition explained in the main text.

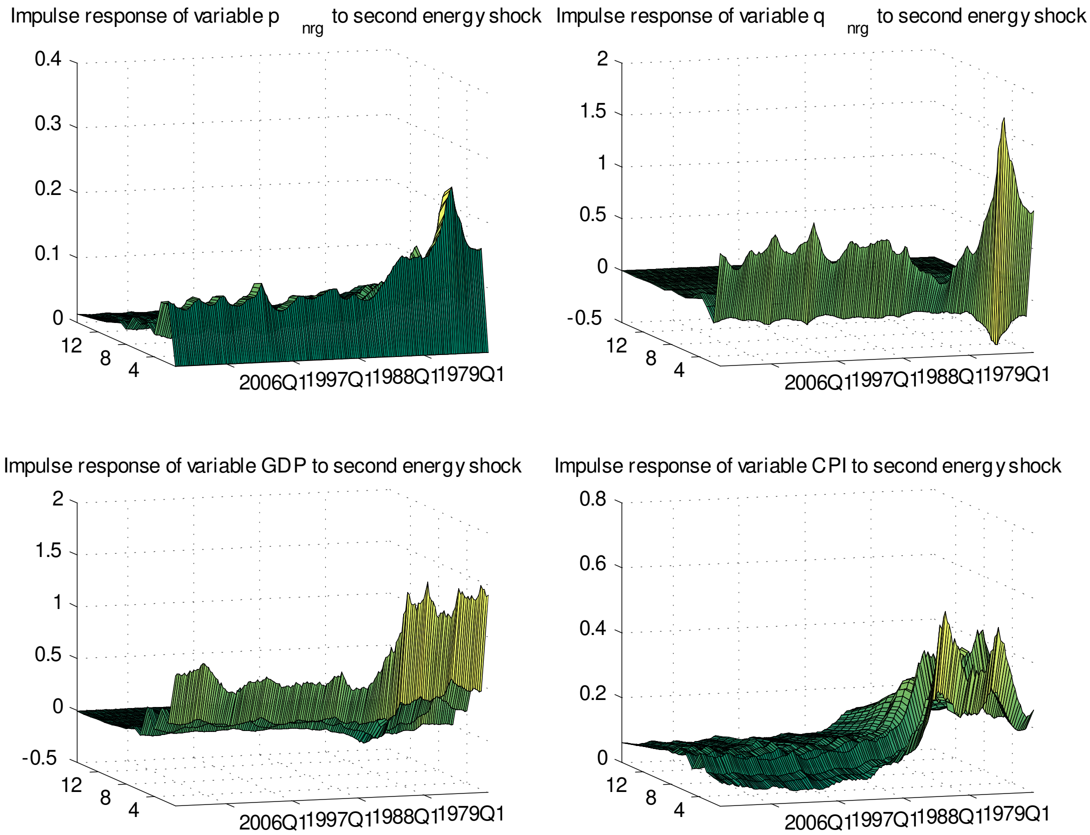

Posterior medians of impulse responses for all 20 horizons and all time-periods (1970Q1-2012Q4). Shock is the second identified ("demand") shock from the decomposition explained in the main text.