

Time Varying SVARs, parameter histories, and the changing impact of oil prices on the US economy

Francesca Rondina[‡]

Institute for Economic Analysis, CSIC and Barcelona GSE

This version: September 2012

comments welcome

Abstract

This work studies the changes in the relationships between oil prices and US domestic variables in the postwar period using a time-varying VAR model of the economy. The paper proposes a new framework, which departs from most of the previous literature with respect to the assumptions on the evolution of the covariance matrix of the reduced form vector of innovations. The VAR model resulting from these assumptions features multivariate stochastic volatility that follows a Wishart autoregressive process. Relative to the standard procedure adopted in the literature, this framework is able to capture a larger variety of time-varying features of the data, and provides some additional insights on the intertemporal dependence between the parameters of the VAR. In the environment under analysis, I find that in specific periods, in particular in the early/mid 1980s, changes in the volatility of the model variables are due to temporary shifts in their contemporaneous relationships rather than to changes in the variance of the exogenous shocks. In addition, the covariances between the residuals to oil inflation and US core inflation and oil inflation and the US growth rate exhibit a strong intertemporal dependence, which might signal some degree of endogeneity of oil prices to US domestic variables.

Keywords: time-varying VARs, parameter histories, oil price shocks, pass-through.

JEL classification: C11, E47, E52, Q43

*Contacts: Campus UAB, 08193 Bellaterra, Barcelona, Spain; tel: +34 935 806 612, fax: +34 935 801 452, email: francesca.rondina@iae.csic.es

[‡]I would like to thank Luca Gambetti, Albert Marcet, Giorgio Primiceri and Juan Rubio-Ramirez, for helpful discussions and suggestions. I have also benefited from the comments received from participants at the CEF conference in Prague and RCEF conference in Toronto. All errors remain my own. Financial support from the Government of Catalonia and the Spanish Ministry for Science and Innovation (Programa JAE-Doc, Junta para la Ampliación de Estudios co-financed by the FSE, and Programa Operativo ECO2011-28348) is gratefully acknowledged.

1 Introduction

This paper investigates the changes in the relationship between oil prices and US domestic variables in the postwar period. In particular, this work focuses on the analysis of the time-varying impact of oil prices on core CPI inflation and GDP growth, and studies the role of monetary policy in these changes. The empirical examination is performed in the context of a VAR framework in which the coefficients of the model and the covariance matrix of the reduced form innovations are allowed to change over time. While this type of framework has already been employed in a number of contributions in the empirical literature, the present paper departs from these previous studies with respect to the specific assumptions on the time-variation of the parameters of the model. This work discusses the central features of this new approach and examines the additional information that can be obtained from its implementation in the analysis of the impact of oil prices on the US economy, thus extending and integrating the results of previous research in this area.

In recent years, a number of contributions in the macroeconometric literature have investigated the development and implementation of techniques that can account for the possibility of structural changes in the true data generating process of the economy. Among these contributions, an important role is covered by the studies that have focused on the definition and estimation of VAR frameworks that incorporate time variation in the relationships between the variables of the model. Significant progresses have been made in this area, different methods for the implementation of this type of analysis have been introduced, and time-varying VAR models have been adopted, and are increasingly employed, to investigate economic relationships in a number of different environments. This work proposes a new approach for the estimation of VAR models with time-varying coefficients and covariance matrix. The main departure of this approach from the previous literature resides on the definition of the time variation of the parameters, and in particular on the assumptions about the evolution of the covariance matrix of the reduced-form innovations. These assumptions have important consequences on the ability of the model to capture some specific the time-varying features of the data, in particular the presence of temporary shifts in the elements of the reduced-form covariance matrix. In addition, the framework developed in this paper also provides some relevant information about the intertemporal relationships between the variances and covariances of the residuals in the reduced-form VAR model, which can be employed to expand the analysis

and improve the understanding of the historical patterns of the variables of interest.

This paper offers a methodological contribution to the literature on time-varying VARs, and in particular to the branch originating from the seminal work of Cogley and Sargent (2001, 2005) and Primiceri (2005). The approach initially proposed by Cogley and Sargent (2001) modeled the time changes in the relationships between the endogenous variables in the VAR in terms of variations in the coefficients of the model. Subsequently, Cogley and Sargent (2005) introduced time variation in the variances of the innovations, and finally Primiceri (2005) extended this framework to include changes in the covariances between the innovations of the model. As in the works of Cogley and Sargent (2001, 2005) and Primiceri (2005), this paper assumes time variation both in the coefficients and in the covariance matrix of the innovations of the VAR framework, and uses Bayesian methods for estimation purposes. The approach that I proposed in this study, however, departs from these previous contributions in terms of the way in which the covariance matrix is assumed to change over time. The Cogley-Sargent-Primiceri method is developed on the assumption of a random walk time variation of the covariances and log standard deviations of the innovations in a structural VAR obtained imposing a specific ordering to the variables in the model. The approach employed in this work, on the other hand, models the evolution of the parameters directly in the reduced-form covariance matrix. In particular, the specific law of motion that I adopt is taken from the finance literature on multivariate stochastic volatility models, in particular from Philipov and Glickman (2006) and Rinnergschwentner et al. (2011), and it assumes that the inverse of this covariance matrix is a Wishart distributed random variable with time-varying scale matrix.

The framework that I use in this paper exhibit a number of relevant features that have important implications for the underlying properties of the time variation of the parameters, and for the interpretation of the intertemporal relationships between them. First, the approach presented in this study departs from the random walk assumption in the time variation of the elements of the reduced-form covariance matrix, and introduces a persistence coefficient which is estimated together with the other parameters. This feature allows the VAR model to capture some characteristics of the data that are not well represented by a random walk process, as for instance the presence of temporary shifts in the variances and covariances of the reduced-form residuals. Second, the way in which the time variation is modeled in this work implies some clear intertemporal dependence between the parameters of the VAR, which can be used to obtain additional insights on

the relationships between the variables of interest. I believe that it is also important to emphasize that the approach proposed in this work does not preclude the type of investigation that is usually developed based on the Cogley-Sargent-Primiceri method. Thus, impulse-response analysis, variance decomposition, and many other commonly performed exercises can still be implemented in the framework introduced in this paper. However, for the reasons that I just explained, I believe that the procedure employed in this study can provide a better representation of the data and offer additional information that can be used in the interpretation of the results of the analysis. These advantages, however, come at the cost of a slightly more complex estimation procedure, and a longer computational time.

This work adds to the time-varying VAR literature by introducing a new approach for modeling parameter changes in a VAR model and by detailing the Bayesian techniques that can be employed for the implementation of this method. In addition, this paper proposes an application of this approach to the analysis of the time-changing impact of oil prices on the US economy, thus also offering a contribution to the empirical literature in this area. In recent years, several studies have opted for the use of time-varying VAR models to investigate the causes of the postulated changes in the effects of oil price shocks on a number of economic variables over time (most notably, Baumeister and Peersman, 2009, and Clark and Terry, 2010). I present an extension of these previous works, which focuses on the analysis of the changes in the impact of oil prices on core CPI inflation and GDP growth, and in the intertemporal and contemporaneous correlations implied by the model.

The results obtained from the empirical analysis are important. The ability of the framework implemented in this paper to capture a larger variety of time-varying features of the data implies that the interpretation of some specific events in the postwar history departs in a relevant manner from the results emerging from previous studies focusing on similar environments. A striking example is provided by the early/mid 1980s, in which I find that changes in the volatility of the model variables are due to temporary shifts in their contemporaneous relationships rather than to changes in the variance of the exogenous shocks. I also find that the covariances between the residuals to oil inflation and US core inflation and oil inflation and the US growth rate exhibit a strong intertemporal dependence, which might signal some degree of endogeneity of oil prices to US domestic variables. In terms of the pass-through, my results support the evidence of a more moderate response of core inflation to oil price shocks since the mid 1980s. However, even before this

date, the model identified using sign restrictions attributes a smaller role to oil price shocks that are exogenous to the US economy compared to the previous works using time-varying VARs estimated using ordering restrictions. Finally, in terms of policy response, the analysis of the intertemporal relationships in the impact of shocks on the variables of interest seems to suggest that the policy reaction to the effects of oil price shocks on oil inflation was mostly related to specific events in the oil market, while the response to the pass-through to core inflation appears to be more systematic over time.

The remainder of the paper is organized as follows. Section 2 describes the class of models under analysis, specifies the assumptions about the time variation of their parameters and delineates the differences between the approach proposed in this work and the standard procedure adopted in the literature. Section 3 provides the details for the empirical implementation of this approach, and section 4 describes its application to the study of the time varying impact of oil prices on the US economy. Section 5 concludes.

2 The model

I am interested in studying the characteristics of a class of structural VARs with time-varying coefficients, covariances and variances of the vector of innovations. Using the same notation as in Primiceri (2005), I start by considering the following reduced-form model:

$$y_t = c_t + B_{1,t}y_{t-1} + \dots + B_{p,t}y_{t-p} + u_t \quad t = 1, \dots, T.$$

where y_t is a $n \times 1$ vector of endogenous variables, c_t is a $n \times 1$ vector of intercepts, $B_{1,t}, \dots, B_{p,t}$ are $n \times p$ matrices of coefficients, and u_t is a vector of innovations with time-varying covariance matrix: $E(u_t u_t') = \Omega_t$. This model can be rewritten equivalently as:

$$y_t = X_t' B_t + u_t \tag{1}$$

where $X_t' = I_n \otimes [1, y_{t-1}', \dots, y_{t-p}']$ and B_t is a $n(1 + np) \times 1$ vector that incorporates all the coefficients in c_t , and $B_{1,t}, \dots, B_{p,t}$. Let B^T and Ω^T denote histories of the parameters up to time T .

The model described by (1) can be interpreted as the reduced-form representation of the struc-

tural VAR:

$$y_t = X_t' B_t + A_t \varepsilon_t \quad (2)$$

where ε_t is a vector of structural shocks of the economy, with $E(\varepsilon_t \varepsilon_t') = I$, and the matrix A_t satisfies: $A_t A_t' = \Omega_t$, $t = 1, \dots, T$. In this framework, A_t^{-1} defines the time t contemporaneous relationships between the endogenous variables, and the time t structural coefficients can be obtained as $A_t^{-1} c_t$ and $A_t^{-1} B_{i,t}$, $i = 1, \dots, p$. Thus, the histories B^T and A^T can be used to characterize the time-varying structural relations between variables implied by model (2).

The empirical literature on time-varying VARs is typically interested in studying the information arising from the patterns of the structural parameters of the model in the specific environments under analysis. The techniques that have been developed for this purpose are based on Bayesian estimation of the joint posterior probability of the parameters of interest, given the available data. The standard procedure, originating from the work of Cogley and Sargent (2001, 2005) and Primiceri (2005), is based on modeling the time-variation of the parameters in the vector B_t , and in the elements of the matrices C_t and Σ_t , which are obtained from the Cholesky decomposition of Ω_t : $\Omega_t = C_t^{-1} \Sigma_t \Sigma_t' C_t^{-1}$. The matrix C_t is lower triangular, with ones in the main diagonal, while Σ_t is a diagonal matrix. The standard assumption is that the elements of B_t , the non-zero and non-one elements of C_t , and the log of the diagonal elements of Σ_t , evolve according to a driftless random walk process in which the innovations are Gaussian random variables. Given this assumption, it is possible to derive the conditional posteriors of the parameters of interest, and a MCMC algorithm, more specifically a Gibbs sampling procedure, can be used to obtain draws from the joint posterior of $(B^T, C^T, \Sigma^T, \tilde{V})$, where \tilde{V} is a matrix of hyperparameters of the model (see Primiceri, 2005, for more details).

For reasons that I will explain in more detail in the next section, I depart from the previous literature and model the time-variation directly in the reduced-form vector of coefficients B_t and covariance matrix Ω_t . More specifically, while for B_t I adopt the same approach as in Cogley and Sargent (2001, 2005) and Primiceri (2005), for the covariance matrix Ω_t I borrow from the finance literature, and in particular from the contributions of Philipov and Glickman (2006) and Rinnergschwentner et al. (2011), and assume that Ω_t^{-1} has a Wishart distribution with time-varying scale matrix. Thus, the dynamics of the model's parameters are described by the following

expressions:

$$B_t = B_{t-1} + v_t \quad (3)$$

$$\Omega_t^{-1} \mid k, S_{t-1} \sim Wish(k, S_{t-1}) \quad (4)$$

$$S_t = 1/k \left(G^{1/2} \right) \left(\Omega_t^{-1} \right)^d \left(G^{1/2} \right)' \quad (5)$$

where v_t is a vector of innovations with $N(0, Q)$ distribution, G is a positive definite symmetric matrix (and $G^{1/2}$ denotes the lower triangular matrix obtained from its Cholesky decomposition), k are the degrees of freedom in the Wishart distribution, and d is a scalar. The quadratic form of S_t ensures that the covariance matrices are symmetric positive definite. As remarked in Philipov and Glickman (2006), the matrix G and the parameter d play an important role in the dynamic behavior of the covariance matrix Ω_t . Notice that in this setup the conditional expectation of Ω_t^{-1} can be written as:

$$E \left(\Omega_t^{-1} \mid G, \Omega_{t-1} \right) = \left(G^{1/2} \right) \left(\Omega_{t-1}^{-1} \right)^d \left(G^{1/2} \right)' \quad (6)$$

This expression highlights how the matrix G (or more precisely the inverse of the matrix G) provides information on the intertemporal dependence between the elements of Ω_t , and on the relative importance of the other variances and covariances on the pattern of each of these elements. The next section will discuss how this feature can be used for inference. The parameter d , on the other hand, is a measure of persistence in the intertemporal relationship between the elements of Ω_t . Values of $|d| > 1$ imply nonstationary dynamics of Ω_t , while values in the interval $[-1, 0)$ generate dynamics in which Ω_t alternates between powers of inverses.¹ For this reason, as in Philipov and Glickman (2006) and Rinnergschwentner et al. (2011), I will restrict my analysis to values of d in the interval $[0, 1]$. Notice that the extremes of this interval represent some interesting special cases. If $d = 0$, the covariance matrix Ω_t exhibit no time-variation:

$$E \left(\Omega_t^{-1} \mid G, \Omega_{t-1} \right) = G$$

while if $d = 1$, the value of the covariance matrix at time $t - 1$ is fully reflected in its current value, with relationships between elements dictated by the matrix G . One final remark needs to

¹See Philipov and Glickman (2006) for a more detailed discussion.

be made with respect to k , the degrees of freedom of the Wishart distribution of Ω_t^{-1} . In order for Ω_t^{-1} to be invertible with probability one, we need k to be larger than the dimension of Ω_t^{-1} ; For this reason, k will be restricted to assume only values larger than n . When the matrix G is symmetric positive definite, the parameter d is included between 0 and 1, and k is greater than n , then the autoregressive stochastic matrix process for the covariances is well defined (see Philipov and Glickman, 2006, and Rinnergschwentner et al., 2011).

As in the previous literature, I will estimate the time-varying VAR model described by (1) and (3)-(5) using Bayesian techniques. Given some assumptions on the prior distributions of the parameters of interest, the conditional posteriors can be analytically derived, and a Metropolis-Hasting algorithm can be implemented to obtain draws from the joint posterior of (B^T, Ω^T, Q, V) (where for simplicity of notation, V encloses the parameters of the Wishart distribution for Ω_t^{-1} , i.e. G, k and d).

2.1 Main properties of the model

This section provides a more detailed description of the main characteristics of the model used in this paper, and discusses the additional information that the assumptions on the time-variation of Ω_t adopted here might convey. This information can provide a better understanding of the patterns of the parameters of interest in the period under analysis, and can be employed to improve their interpretation in the context of the class of structural VAR models described by (2).

First, I think that it is important to remark that the method proposed in this work still provides all the information that can be obtained from the standard approach adopted in the literature, so that all the common exercises and counterfactuals that have been developed based on the Cogley-Sargent-Primiceri technique can still be performed here. As previously mentioned, the standard procedure employed to estimate time varying VARs delivers draws from the posterior: $p(B^T, C^T, \Sigma^T, \tilde{V} | y^T)$, where y^T is the history of the endogenous variables of the model up to time T . The vector B_t is the vector of reduced-form coefficients of the model, and the matrices C_t and Σ_t are obtained from the Cholesky decomposition of the covariance matrix Ω_t . Notice that while this specific decomposition is chosen mainly because it considerably simplifies the analysis of the time variation and the estimation of the model, its adoption has some important implications. In particular, this setup implies that the order of the variables in the VAR matters (see Primiceri,

2005, for a discussion). Indeed, this decomposition can be interpreted as corresponding to the particular structural VAR:

$$y_t = X_t' B_t + C_t^{-1} \Sigma_t \varepsilon_t \quad (7)$$

where $C_t^{-1} \Sigma_t$ is a lower triangular matrix. Because of this lower triangular structure, the model in (7) denotes a specific way in which the structural shocks in the vector ε_t affect the variables in y_t , and this implies a specific order in which these variables are determined. In some applications, researchers might be interested in the analysis of time-varying structural VAR models that are not necessarily the one described by (7). The approach that is commonly adopted in this case makes use of orthogonal matrices to define A_t at every point in time starting from C_t and Σ_t :

$$A_t = C_t^{-1} \Sigma_t P_t \quad (8)$$

Because of the properties of orthogonal matrices, we have that:

$$A_t A_t' = C_t^{-1} \Sigma_t P_t P_t' \Sigma_t' C_t'^{-1} = C_t^{-1} \Sigma_t \Sigma_t' C_t'^{-1} = \Omega_t$$

so that any matrix A_t obtained from (8) will give the same Ω_t , which means that all structural VAR models corresponding to different matrices A_t will share the same reduced-form covariance matrix. Usually, sign restrictions are subsequently employed to select a subset of structural VARs that exhibit some required characteristics, and to identify the shocks of interest (see, for instance, Canova and Gambetti, 2009, which is the first contribution that applied this approach in a time-varying environment).

The model that I employ in this work, and the estimation procedure that I will describe in the next section, deliver draws from the joint posterior: $p(B^T, \Omega^T, Q, V | y^T)$ rather than $p(B^T, C^T, \Sigma^T, \tilde{V} | y^T)$. However, the histories (C^T, Σ^T) can easily be obtained from the Cholesky decomposition of the draws of Ω^T generated using this technique, and patterns of the structural matrix A_t can subsequently be obtained using the procedure described by (8). It follows that the analysis that is usually carried on using the standard approach, for instance identification and inference performed using sign restrictions on the elements of A_t in each period t , can still be performed using the method proposed in this paper.

While, the approach adopted in this paper can be used to obtain the same information as the procedure based on the estimation of the histories of C_t and Σ_t , I believe that there are some characteristics that would make it preferable in a number of environments. In particular, some of the parameters in (4) and (5) define important properties of the time variation of Ω_t , and might convey relevant information which can be used in the study and interpretation of the patterns of the variables of interest in the period under analysis.

The main feature in which the stochastic multivariate volatility model described by (3)-(5) departs from the time-varying volatility frameworks studied by Cogley and Sargent (2005) and Primiceri (2005) is in the introduction of the parameter d , which measure the intertemporal persistence in the process for Ω_t . Because of the random walk assumption on the time variation of the non-zero and non-one elements of C_t , and the log of the diagonal elements of Σ_t , the standard model employed in the literature implies a value of d equal to one.² Although it is true that Primiceri (2005) discusses an extension of the baseline framework that allows the coefficients to follow a more general AR process, in practice the vast majority of the empirical work employing time-varying VARs adopts the assumptions of a random walk. In the approach proposed in this paper, the value of d is estimated together with the other parameters of the model, and can therefore take values that are lower than one. I believe that this is an important feature of the framework used here, because it allows the stochastic covariance matrix Ω_t to reflect temporary shifts in the reduced-form variances and covariances. This property might be very important in VAR models that incorporate one or more variables exhibiting periods of higher volatility followed by others of relative more stability, because in this case the framework proposed in this paper would be able to capture and describe the patterns of interest in a better way. The empirical application described in section 4 is an example of such an environment. Indeed, the history of oil prices is clearly characterized by alternation of periods of higher and lower volatility. For this reason, I think that the possibility to account for temporary fluctuations in the elements of Ω_t makes the approach proposed in this paper more adapt to the analysis of the impact of oil price changes on the US economy than the standard procedure used in the literature.

²This statement has been verified numerically. More specifically, I generated an history Ω^T using the laws of motion adopted in Primiceri (2005), and then I estimated the model described by (4) and (5). The draws of d that I obtained from this procedure were all included between 0.95 and 1. I repeated the exercise using different histories Ω^T generated under alternative assumptions on the parameters of the laws of motion for the relevant elements of C_t and the log of the diagonal elements of Σ_t , and I obtained the same result.

A second merit of the model described by (4) and (5) lies in the information provided by the matrix G . As previously discussed, the inverse of the matrix G gives a measure of the intertemporal dependence between the elements of Ω_t . In addition to being useful for the analysis and interpretation of the results, the information delivered by G can also be very important for the choice of the identification assumptions that need to be implemented in order to obtain the histories A^T from the estimated histories Ω^T . A large fraction of the studies employing time-varying VAR models make use of ordering restrictions to identify the structural framework of interest for the analysis. As an example, in the environment examined in the empirical section of the paper, the vast majority of previous works adopt this type of restrictions, and the common procedure is to assume that oil prices are determined before all the domestic variables included in the VAR model.³ I believe that the matrix G can offer relevant information on whether a given restriction on the order in which the variables are determined is supported or opposed by the data. For instance, it would be difficult to argue that a given variable is exogenous relative to the others if the elements of G^{-1} suggest that the variance of its reduced-form innovations is highly related to the covariances with the other innovations. On the other hand, near zero values of the off-diagonal elements of G^{-1} imply that the variance of the reduced-form innovations to a given variable depends only on its past values, which would support the assumption that this variable is determined before the others. This same type of reasoning can also be employed to decide whether ordering assumptions are suitable to identify the model of interest in the environment under study, or whether alternative approaches (sign restrictions, long run restrictions,...) are necessary.

Another interesting property of the stochastic volatility framework proposed in this paper is that it defines a clear and intertemporal relationship between the elements of A_t and, for this reason, it promotes the study of the intertemporal changes in the parameters of the model and of their impact on the relationships between the variables of interest. For instance, a natural measure that could be examined is the distribution of the changes in the response of the variables of interest to some selected shock(s), and the correlations of these changes over time. In addition, this framework also emphasizes the determinants of the time relationships between the elements of A_t , and the reasons why these might have changed over time. Given the orthogonal matrices P_t and P_{t-1} , the intertemporal dependence between the elements of Ω_t defines a relationship between the elements

³See section 4 for a more detailed discussion.

of A_t and A_{t-1} , which can be used in the analysis. The matrix G can also be interpreted as determining mean reversion properties of the covariance matrix Ω_t , which will also be reflected on the matrices A_t . Notice that since both G and d are constant over time, this framework implies that the causal relationships between elements of Ω_t do not change across periods. This will generally not be the case for the intertemporal dependence between the elements of A_t , because this will also be determined by the rotation matrices used to define A_t at each t , and by the restrictions imposed to identify the histories A^T . I believe that in a number of economic environments it might be very interesting to investigate how the relationships between elements of A_t and A_{t-1} have evolved over time, and the role that the different parameters of the model played in these changes.

3 Empirical Implementation

In this section, I provide further information about the implementation of the approach proposed in this paper. The procedure is developed in two different stages. In the first one, I employ a MCMC algorithm to estimate the model described by (1) and (3)-(5) and obtain draws from the posterior distribution of interest: $p(B^T, \Omega^T, Q, V | Y^T)$. In the second, I use these draws to define possible histories A^T by employing (8) given rotation matrices P_t for each draw and each time t . The second stage of the analysis is quite standard in the literature, and therefore it will be discussed only briefly here. The first stage, on the other hand, includes the main methodological innovations of this work, and for this reason it will be illustrated in more detail. A more technical description of the procedure is given in Appendix 1; in this section I will explain the assumptions on the prior distributions for the parameters of interest and the derivation of the relevant posteriors.

The set of parameters of interest can be separated into two blocks, one including the history of the vector of reduced form coefficients B_t and its hyperparameter Q , and the second incorporating the history of the covariance matrix Ω_t and the parameters in V . The prior distributions are assumed to be such that:

$$p(B^T, \Omega^T, Q, V) = p(B^T) p(Q) p(\Omega^T | V) p(V)$$

The choice of these distributions follows the literature. For the first block, the assumptions are

as in Cogley and Sargent (2001, 2005) and Primiceri (2005). In more detail, a normal prior is chosen for B_0 , and an inverse-Wishart for Q .⁴ The conditional posteriors $p(B^T | \Omega^T, Q, V, Y^T)$ and $p(Q | B^T, \Omega^T, V, Y^T)$ can be derived using these prior distributions and the assumptions on the time-variation of B_t . The conditional posterior of B^T is a product of Gaussian densities, and a draw from it can be obtained using a simulation smoother as the one described in Carter and Khon (1994). The matrix Q has an inverse-Wishart conditional posterior, and its value can be drawn directly from this distribution.

For the second block of parameters, I follow Rinnergschwentner et al. (2011). The prior for the inverse covariance matrices is defined as a Wishart distribution, given the parameters in V and consistently with (4). The elements of V are assumed to have independent prior distributions. For G , the prior is defined on G^{-1} , and it is assumed to be a Wishart distribution. For k and d , uniform distributions are chosen. Given these assumptions on the prior distributions and the time variation of Ω_t described by (4)-(5), the analytical formulas for the conditional posteriors can be obtained, although they are complicated expressions and do not correspond to known distributions (see Rinnergschwentner et al., 2011, for the complete derivation).

The MCMC algorithm works by sampling values from the conditional posterior distributions of the parameters of interest. After an initial burn in period, subsequent draws will correspond to draws from the joint posterior $p(B^T, \Omega^T, Q, V | Y^T)$. In the specific setup used in this paper, the approach that needs to be employed to sample from the conditional posteriors is different for the first and second block of parameters. For B^T and Q , the conditional posteriors are well known distributions. This implies that we can draw directly from them, using a Gibbs sampler type of algorithm. On the other hand, the form of the conditional posteriors for Ω^T and the parameters in V is non-standard. For this reason, a Metropolis-Hasting algorithm is necessary to obtain draws from these posteriors. This algorithm works by obtaining proposed draws of the variables of interest from a known distribution, and by accepting or rejecting these proposed draws based on an acceptance rate computed using the expressions for the conditional posteriors. This procedure is more involved than the Gibbs sampler, and typically requires longer burn-in periods for the draws to converge to the joint posterior of interest.⁵ However, in the exercise that I will present in the

⁴The "virtual" prior for B^T can then be obtained using the prior for B_0 and the assumption on the time variation of B_t , see Cogley and Sargent (2001).

⁵Actually, the Gibbs sampler is a special case of the Metropolis-Hasting algorithm, in which the draws are generated

next section, I found that the efficiency of the MCMC procedure is greatly improved if J iterations of the Metropolis-Hasting step are performed for each draw of B^T and Q obtained from the Gibbs sampler step.⁶

The draws of (B^T, Ω^T, Q, V) obtained from the MCMC algorithm, can finally be used to generate a set of possible histories A^T . The approach that I employ is the following. First, I randomly draw one history $\Omega^{T(i)}$ from the set obtained from the MCMC algorithm. Then, I compute the Cholesky decomposition of these covariance matrices at each point in time: $\tilde{C}_t^{(i)} \tilde{C}_t^{(i)'} = \Omega_t^{(i)}$, $i = 1, \dots, I$. Finally, for each t I draw an orthogonal matrix P_t , and set: $A_t^{(i)} = \tilde{C}_t^{(i)} P_t^{(i)}$.⁷ The value $A_t^{(i)}$ can be kept or discarded based on the specific assumptions relative to the structural model of interest to the researcher. The algorithm can be stopped once the desired number of histories A^T is reached.

4 The time-varying impact of oil prices on the US economy

In this section, I use the proposed time-varying model to study the time changes in the impact of oil prices on the US economy. The economic literature has long been interested in the analysis of the effects of oil price shocks on a number of economic variables. Many previous contributions in this area have made use of VAR models to study the relationships between oil prices and a set of variables of interest, and to investigate possible changes in these relationships over time. For instance, Bernanke, Gertler and Watson (1997, 2004), Blanchard and Gali (2007), Hamilton and Herrera (2004), Herrera and Pesavento (2009) and Kilian and Lewis (2009) use VAR frameworks to discuss the role of monetary policy in the downturns following the large oil price shocks of the postwar period, and to investigate the extent to which the Fed's response to changes in oil prices contributed to their allegedly reduced impact on several economic variables since the mid 1980s. Lippi and Nobili (2008) use a structural VAR with constant parameters to study the effects of oil

directly from the known conditional distributions.

⁶This approach is justified by the fact that the Metropolis-Hasting algorithm can be performed one block or one variable at the time. See Chib and Greenberg (1995) and Hasting (1970) for a discussion. The number J , however, needs to be small, otherwise the gains obtained from the faster convergence of the algorithm are lost because of the greater autocorrelation of the draws. In empirical applications, I found that values of J equal to 3 – 5 work well in general.

⁷The orthogonal matrix P_t can be obtained by drawing an arbitrary $n \times n$ independent standard normal matrix X , computing the QR decomposition of X , and setting $P_t = Q$. See Rubio-Ramirez et al. (2010) for a discussion of this approach.

supply and demand, as well as shocks in the world economy, on US real activity and inflation, while Baumeister and Peersman (2008) employ a time-varying VAR approach to perform a similar analysis. Finally, Clark and Terry (2010) implement the Cogley-Sargent-Primiceri approach to examine the pass-through of energy price inflation to US core inflation.

Because an extensive analysis of the time-varying impact of oil prices of a number of economic variables has already been provided in the mentioned literature, this section will mainly investigate the differences and the additional information that can be obtained using the approach proposed in this paper. More specifically, I will focus on the estimated values of the parameter d and matrix G and I will discuss their consequences for the results of the analysis and for the interpretation of the relationships between the variables of interest. In addition, I will study the intertemporal relationships between the parameters in a number of possible structural VAR models of the economy identified using alternative restrictions on the impact of selected shocks on the endogenous variables of the model.

The VAR model includes 4 variables: oil price changes, core inflation, a measure of US real activity, and the nominal interest rate.⁸ The data is quarterly, and the estimation period goes from 1970 – I to 2010 – IV (data from 1957 – II to 1969 – IV is used to compute the parameters of the prior distribution required for the Bayesian estimation procedure). The VAR model is specified with 3 lags of the variables of interest, as in Clark and Terry (2010). This is a longer number of lags compared to those used in some of the previous literature (for instance, Lippi and Nobili, 2008), and it was chosen to account for the fact that, as pointed out by Hamilton and Herrera (2004), important effects of oil prices on the US economy might happen at longer lags.

The prior distributions were specified as described in the previous section. The parameters of these distributions were chosen as follows. In the normal prior for B_0 , the mean was set as the OLS estimate of a model with constant parameters obtained using a training sample that includes the available data up to 1969 – IV. The variance was set as the variance of this OLS estimate, multiplied by 4 in order to increase the dispersion of the prior and make it less informative. The inverse Wishart prior distribution for Q was assumed to have $k_{Q_0} \times (n_B + 1)$ degrees of freedom,

⁸The data used in this empirical application is as follows. Oil price changes (or oil price inflation) are measured as the changes in the nominal price of oil, expressed in log terms. The nominal price of oil is the West Texas Intermediate spot oil price. Core inflation is the annualized difference in log core CPI, where core CPI is the “CPI for all urban consumers: all items less food and energy”. The measure of real activity is the growth rate, computed as the log difference in real GDP. Finally, the last variable included in the VAR is the annualized Federal funds rate.

and scale parameter equal to the variance of the OLS estimate for B_0 multiplied by $k_{Q_0} \times (n_B + 1)$. The value of k_{Q_0} was set equal to $(0.005)^2$, and n_B is the length of the vector B_t , i.e. $n(1 + np) \times 1$. For Ω^T , the initial value Ω_0 was obtained from the same OLS estimates of a constant model used for B_0 . For the other parameters in the Wishart distribution of Ω_t^{-1} , the initial values were set as: $G_0 = I$, $k_0 = n + 1$, $d_0 = 0.5$. Finally, the parameters of the proposal distributions in the Metropolis-Hasting procedure were assumed to be: $\sigma_{G^{-1}}^2 = 0.005^2$, $\sigma_k^2 = 2.5^2$, $\sigma_d^2 = 0.01^2$ (see the Appendix for more details about the Metropolis-Hasting procedure).

I performed 350,000 iterations of the algorithm, discarding the first 50,000 as burn-in period. I saved one every 300 of the remaining iterations in order to break the autocorrelation of the draws. The total number of draws of $(B^T, \Omega^T, Q, G, k, d)$ that I kept was thus of 1,000. I selected a value of $J = 5$, meaning that 5 iterations of the Metropolis-Hasting procedure for (Ω^T, G, k, d) were made for each draw of (B^T, Q) obtained from the Gibbs sampler. The acceptance rates in the Metropolis-Hasting procedure were all between 30% and 50%, with the exception of the covariance matrices Ω_t for which the rates were lower. However, the autocorrelations between draws of Ω_t were also very small, signaling that the lower acceptance rates were not an issue for the convergence of the algorithm. I checked the 20th order autocorrelations across draws for most of the estimated parameters and, with some few exceptions, the values were all very small. In general, the algorithm seems to perform well and to converged quickly (after about 8,000 iterations).

Table 1 reports the medians, together with the 16th and 84th percentile values, of k , d , and the diagonal elements of G^{-1} in the retained 1,000 draws. The table also shows the acceptance rates of the Metropolis-Hasting algorithm for these parameters. The estimated value of d suggests some high persistence in the intertemporal relationships between the elements of Ω_t . The diagonal elements of G^{-1} , on the other hand, provide information about the relationship of each reduced-form variance at time t with its own value at $t - 1$. The variables in the VAR are in the order: oil price changes, core inflation, growth rate, Federal funds rate. Thus, table 1 shows that the volatility of the reduced-form innovations to oil price changes is much more dependent on its own past values than the volatilities of the other reduced-form innovations. Table 2 provides an even better picture of the intertemporal relationships between the variances and covariances in Ω_t by reporting the median values of the elements of G^{-1} .

This table suggests that the off diagonal dependence between the innovations to the VAR

variables is quite small, with the relevant exception of oil price changes. Two important observations can be made from these results. First, even if it is true that the value of d is quite high, it is still significantly different from one, the value implied by the previous approaches used to model time-varying VARs. It follows that in the environment under study the time-variation of Ω_t is better described by shifts that are quite persistent, but not permanent. This result was somehow expected since, as mentioned above, the histories of some of the endogenous variables in the VAR model, in particular oil prices, are characterized by periods of higher volatility followed by more stable times, and these patterns cannot be captured by a random walk process. I will discuss below as this value of d smaller than one has important consequences for the relationships between the variables of interest in specific time periods. The second observation that we can make refers to the data reported in table 2. This table clearly shows that there is some intertemporal dependence between the volatilities and covariances of the innovations to oil inflation and US core inflation and growth rate. These relationships suggests that at least some component of oil price changes might actually be endogenous to the US economy and that, for this reason, the ordering restrictions used in most of the previous studies cited above might actually not be the best choice in this environment.

In order to investigate in more detail the extent to which the approach proposed in this paper is able to capture the time-varying features of the data, it is useful to compare the results obtained here with those reported by Clark and Terry (2010). This previous work implements the Cogley-Sargent-Primiceri procedure to study the changes in the impact of oil prices on the US economy, in a VAR framework in which oil prices are assumed to be exogenous to the domestic variables included in the model. The data used by Clark and Terry (2010) is very similar to the data that I employed in the present paper, except for the measure of real activity for which I use the growth rate while they adopt alternative variables but report their results only for the unemployment gap. Figure 1 reproduces the time-varying standard deviations of the reduced-form residuals, and the standard deviations of the structural shocks identified using ordering restrictions, with oil inflation ordered first in the VAR model. It can be verified that the patterns in these figures are similar to those reported by Clark and Terry (2010) (see figure 4 in this work). The standard deviations of the innovations to oil inflation and core inflation seem to be more volatile under the approach adopted here; this is most probably a consequence of the value of the parameter d which, as previously discussed, is lower than one. However, the most striking difference between the results in Clark

and Terry (2010) and those emerging from the framework proposed in this paper is in the pattern of the standard deviation of the monetary policy shocks, in particular in the early 1980s. In Clark and Terry (2010), the standard deviation of the structural monetary policy shock exhibits the same pattern as the standard deviation of the innovations to the Federal funds rate obtained from the reduced-form model. This is clearly not the case for the results reported in figure 1, in which in several occasions the standard deviations of the structural shocks and reduced form innovation have different patterns. From a technical point of view, these differences are due to the presence of the parameter d , which allows the stochastic covariance matrix Ω_t to reflect temporary shifts in the reduced-form variances and covariances. The consequences for the interpretation of the results of the analysis are of primary importance. Indeed, the conclusion that we can derive is that if we adopt a model in which the time variation of the covariance matrix Ω_t is able to incorporate changes that are nonpermanent, then some of the shifts in the variances and covariances of the reduced-form innovations might not be attributed to changes in the volatility of the exogenous structural shocks, but rather to changes in the contemporaneous response of the endogenous variables to these shocks. In the specific environment under analysis, the differences in the patterns of the standard deviations of the structural and reduced-form innovations in the early 1980s are in fact due to changes in the contemporaneous response of the Federal funds rate to the other variables in the model.

While comparing the results obtained using the framework proposed in this paper with those reported in Clark and Terry (2010) can provide relevant information about the consequences of alternative assumptions on the time variation of Ω_t , it is also true that, as previously discusses, the estimated posterior distribution of the matrix G suggests that oil prices might be endogenous to the US economy to some extent. For this reason, I decided to depart from the ordering identification assumptions adopted in Clark and Terry (2010) and to perform the second part of the analysis on a structural VAR model identified using sign restrictions. I considered three types of shocks. The first two types of shocks are US domestic shocks. More specifically US demand shocks are assumed to have a positive impact on the growth rate and both oil inflation and core inflation, while US supply shocks are assumed to have a positive effect on the growth rate and oil inflation, but a negative one on core inflation. These shocks will be denoted as $\varepsilon_{d,t}$ and $\varepsilon_{s,t}$ respectively. The third type of shocks are oil shocks, which are assumed to have a positive impact on oil inflation, and a negative one on the growth rate. These shocks will be denoted as $\varepsilon_{o,t}$. Following Lippi and Nobili (2008), the sign

restrictions that I impose are applied at impact only. From the 1,000 draws of $(B^T, \Omega^T, Q, G, k, d)$ that I retained from the MCMC procedure, I originated 5,000 different histories A^T satisfying the restrictions that I just described.

The oil price shocks that I identify here can be interpreted as oil supply shocks, but in fact they are all type of disturbances that increase oil inflation and have a negative impact on US real activity. Thus, an increase in the world demand of oil which increases oil prices in a context in which the US are experiencing a downfall in the growth rate of real GDP, would be interpreted as an oil price shock in this framework. Notice that no assumptions are made on the response of core inflation to oil price shocks, and on the impact of all types of shocks on the Federal funds rate. The patterns of these elements of A_t will be the focus of the empirical study, and for this reason their behavior must be left unrestricted in order to prevent the conclusions of the analysis from being dictated by the chosen assumptions.

Using the set of 5,000 selected histories A^T , I exploit the features of the model employed in this paper and I focus on the analysis of the intertemporal relationships between the responses of the VAR variables to the different shocks. In addition, I investigate three questions that have been previously examined in the literature. The first one refers to the decrease in pass-through of oil inflation to US core inflation, which has been studied by Clark and Terry (2010). Can the approach of this paper confirm this decrease in pass-through? The second and third questions are related to the findings of Kilian and Lewis (2009). Does the approach adopted here corroborate the lack of evidence of a systematic response of monetary policy after 1987, and the finding of a different policy response to different sources of oil price increases prior to this date? In studying these three questions, I will focus in particular on the measures and exercises that are more exemplary of the estimation procedure adopted in this paper.

The preliminary results of the analysis are reported in figures 2 – 8. The first exercise that I performed was in the direction of looking at the response of the endogenous variables to the shocks of interest. Figure 2 reports the impact of $\varepsilon_{o,t}$ on core inflation (panel 1) as well as the fraction of histories in which this impact is positive, i.e. the fraction of histories in which oil price shocks induce an increase in core inflation (panel 2). This figure shows that the number of histories in which the response of core inflation to $\varepsilon_{o,t}$ is positive changes significantly over the period under study. This number is very high for the large oil price increases of the early 1970s, and even if it

decreases for the later years, it still remains quite high for the entire sample period. The magnitude of the response of core inflation to oil price shocks is reported in panel 1, and exhibits a similar pattern. Figure 3 compares the results emerging from the sign restrictions approach adopted in this study with those obtained using a Cholesky decomposition of the draws of Ω^T obtained from the MCMC algorithm. In particular, the figure reports the changes in the impact of oil price shocks on oil inflation and core inflation, and their correlation over time. The model identified using sign restrictions implies a smaller role of oil price shocks in the changes in oil inflation, and smaller variations in the pass-through to core inflation. In all, figures 2 and 3 suggest that, even if the impact of oil price shocks on core inflation became more moderate after the mid 1980s, under the assumptions adopted in this paper the influence of these shocks is more limited during the entire period under analysis relative to the results of previous works using the Cholesky identification assumption (as for instance Clark and Terry, 2010).

Figure 4 shows the response of oil inflation and the Federal funds rate to $\varepsilon_{o,t}$, $\varepsilon_{d,t}$ and $\varepsilon_{s,t}$. Figure 5 reports the same information for all the shocks together, so that the timing of the response to the different shocks can be compared. These figures clearly show how the different types of shocks contributed to oil inflation over time. In particular, the model suggests that US domestic shocks, in particular demand shocks, had an important impact on oil inflation not only in recent years, but also in the 1970s and 1980s. In addition, the bottom panels show that the contemporaneous impact of different types of shocks on the Fed funds rate is very different. Most notably, in the early 1980s the overall response of the Federal funds rate is a composite of positive and negative reactions to the domestic and oil price shocks. Figures 6 and 7 are similar to figures 4 and 5 but report information about the changes in the contemporaneous response to the different shocks instead of the levels. Two more observations can be made from these figures. The first one is that changes the impact of shocks on oil inflation seem to revert back quickly after they happen. Indeed, most of the changes in the top panels of figures 6 and 7 exhibit subsequent changes of the same magnitude and opposite direction shortly afterwards. The second observation refers to the impact of shocks on the Federal funds rate in the early 1980s. This model suggests that the timing of the response to different shocks was not contemporaneous, and that an initial reaction to a demand shock was followed by a positive response to an oil price shock.

Finally, figures 8 focuses on the intertemporal and contemporaneous correlations between the

response of the variables to the different types of shocks. The correlations between the response of oil inflation and the Fed funds rate to all types of shocks is high and positive in periods that correspond to specific events in the oil market, namely large increases in the price of oil. The correlations between the response of core inflation and the Fed funds rate, on the other hand, appear to be less linked to individual events and, with the exception of the recent years, seem to be related to the pattern of the business cycle. In all cases, the contemporaneous relationships between the responses of oil or core inflation and the Fed funds rate are similar for all types of shocks, which can be interpreted as implying that the monetary authorities do not react to the shocks themselves, but rather to their effects on the variables of interest. It is also interesting to compare the right-hand panels of figure 8 with the left-hand panels. The intertemporal correlations between the response of oil or core inflation and the Fed funds rate to shocks are clearly much weaker than the contemporaneous ones. However, while for core inflation there seems to be a pattern in these correlations, which exhibits a similar behavior compared to the contemporaneous ones, this is not the case for oil inflation. Thus, this result shows a more persistent response of the Fed funds to the impact of shocks on core inflation rather than to their effects on oil inflation. This statement is confirmed in figure 9 which reports the intertemporal correlations in the response of the endogenous variables to the different shocks. For core inflation, this correlation is lower for oil shocks relative to the domestic shocks, and the same happens for the Federal funds rate. This implies that these two variables seem to react to oil shocks in a similar way, and in a way that is different from the response to the other types of shocks identified in this exercise. This conclusion could be interpreted as supporting the idea that in the sample period under analysis the policy variable reacted in a more systematic way to the effects of oil shocks on core inflation rather than to their direct impact on oil inflation, and only occasionally responded to specific events in the oil market which caused very large variations in oil inflation.

5 Conclusions

In this paper, I proposed a new approach for the estimation of time-varying VAR frameworks. This technique is developed on assumptions about the time variation of the parameters which refer to the coefficients and covariance matrix of the reduced-form VAR corresponding to the set of structural

VARs of interest. I provided a description of the MCMC procedure that can be implemented to generate draws from the joint posterior distribution of the parameters of the reduced-form model, and I suggested how possible histories of the structural VAR parameters can subsequently be obtained from these draws.

I applied this approach to the analysis of the impact of oil prices on the US economy focusing, in particular, on the information that can be obtained from the parameters of the model and the assumptions on their time variation. I showed that the framework adopted in this paper can provide important insights on the nature of the changes in the relationships between oil price inflation and the domestic variables of interest. I also found that the differences with the results obtained from a framework featuring the standard assumptions on the time variation of the reduced-form covariance matrix might be relevant, particularly in some of the periods under analysis.

I believe that the approach proposed in this paper can be implemented in a number of other environments. As discussed in the main text, in frameworks that include variables characterized by periods of high and low volatility the time variation of the parameters of the model would be better described and characterized using the approach developed in this paper, rather than the standard procedure adopted in the literature. In all, I think that this paper offers a framework that can be employed to obtain relevant information and insights in many time-varying VAR models, and that we could benefit from its further development and application in the future.

Appendix 1

MCMC algorithm to obtain draws from $p(B^T, \Omega^T, V | y^T)$

The Markov Chain Monte Carlo Algorithm used in this work is a combination of the technique developed in Cogley and Sargent (2001) for VAR models with time-varying coefficients, and the technique described in Rinnergschwentner et al. (2011) for multivariate models with stochastic volatility. More specifically, steps 1 – 2 are taken from Cogley and Sargent (2001), while steps 3 – 6 follow Rinnergschwentner et al. (2011). This appendix provides a schematic descriptions of the different steps of the algorithm; for a more detailed discussion see the mentioned works.

Step 1 - Drawing the history of coefficients B^T

Conditional on Ω^T and V , the VAR is linear and has Gaussian innovations with known variance. The density $p(B^T | y^T, \Omega^T, V)$ can be factored as:

$$p(B^T | y^T, \Omega^T, V) = p(B_T | y^T, \Omega^T, V) \prod_{t=1}^{T-1} p(B_t | B_{t+1}, y^t, \Omega^T, V)$$

where

$$B_t | B_{t+1}, Y^t, \Omega^T, V \sim N(B_{t|t+1}, R_{t|t+1})$$

$$B_{t|t+1} = E(B_t | B_{t+1}, y^t, \Omega^T, V)$$

$$R_{t|t+1} = Var(B_t | B_{t+1}, y^t, \Omega^T, V)$$

For each t , $B_{t|t+1}$ and $R_{t|t+1}$ can be drawn using the forward Kalman filter and the backward recursion explained next. Given an initial value for the vector of coefficients and for the variance matrix, the Kalman filter develops a recursion giving values of $B_{t|t-1}$ and $R_{t|t-1}$, and when new information is obtained, $B_{t|t}$ and $R_{t|t}$. The last elements of this recursion are $B_{T|T}$ and $R_{T|T}$. At this point, the draw of $B_{T|T}$ and the output of the filter are used for the first step of the backward algorithm, which provides values of $B_{t-t|t}$ and $R_{t-1|t}$ for any value of t until 0.

More specifically, our model is:

$$y_t = X_t' B_t + u_t$$

$$B_t = B_{t-1} + v_t$$

where the assumption on the measurement error and transition shock is:

$$\begin{bmatrix} u_t \\ v_t \end{bmatrix} \sim i.i.d. N \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \Omega_t & 0 \\ 0 & Q \end{bmatrix} \right) \quad (9)$$

Let:

$$B_{t|s} = E (B_t | y^S, X^S, \Omega^S, Q)$$

$$R_{t|s} = Var (B_t | y^S, X^S, \Omega^S, Q)$$

Then, given $B_{0|0}$ and $R_{0|0}$, a standard Kalman filter delivers:

$$B_{t|t-1} = B_{t-1|t-1}$$

$$R_{t|t-1} = R_{t-1|t-1} + Q$$

$$K_t = R_{t|t-1} X_t' (X_t R_{t|t-1} X_t' + \Omega_t)^{-1}$$

$$B_{t|t} = B_{t|t-1} + K_t (y_t - X_t B_{t|t-1})$$

$$R_{t|t} = R_{t|t-1} - K_t X_t R_{t|t-1}$$

From this recursion, we can obtain the values of $B_{T|T}$ and $R_{T|T}$, which are the mean and variance of the normal distribution used to make a draw for B_T . The draw of B_T and the output of the filter can then be used for the first step of the backward recursion. This recursion will provide $B_{T-1|T}$ and $R_{T-1|T}$, which can then be used to make a draw of B_{T-1} . The backward recursion continues until time zero. For a generic time t , the updating formulas for the backward recursion are:

$$B_{t|t+1} = B_{t|t} + R_{t|t} R_{t+1|t}^{-1} (B_{t+1} - B_{t|t})$$

$$R_{t|t+1} = R_{t|t} - R_{t|t} R_{t+1|t}^{-1} R_{t|t}$$

Step 2 - Drawing the hyperparameter Q

Given that the vector of innovations to the reduced form vector of coefficients B_t is assumed to be uncorrelated with the reduced form innovations u_t , the distribution for Q is independent of the distribution for Ω_t . In addition, given a draw of the history B^T , the innovations v_t are observable and they can be used to define the parameters of the posterior inverse Wishart distribution for Q . A value of Q can then be drawn from this distribution. In more detail, if the prior for Q is an inverse Wishart distribution with scale parameter Q_0 and $k_{Q_0} \times (n_B + 1)$ degrees of freedom, then the posterior will be:

$$Q \mid B^T, y^T \sim IWish(k_{\tilde{Q}}, \tilde{Q})$$

where:

$$k_{\tilde{Q}} = k_{Q_0} \times (n_B + 1) + T$$

$$\tilde{Q} = Q_0 + \sum_{t=1}^T v_t v_t'$$

Step 3 - Drawing the history of reduced form variance-covariance matrix Ω^T

Given the draw of B^T obtained in the previous step, we have:

$$(y_t - X_t' B_t) = u_t$$

where $E(u_t \mid \Omega_t) = N(0, \Omega_t)$.

The value of Ω_t^{-1} can be sampled using an independence chain Metropolis Hasting (*MH*) step, in which the proposal density is $Wish(k, \tilde{S}_{t-1})$ with $\tilde{S}_{t-1} = (S_{t-1}^{-1} + u_t u_t')^{-1}$. The acceptance ratio is:

$$AR = \frac{\left| (\Omega_t^*)^{-1} \right|^{(1-kd)/2}}{\left| (\Omega_t^{[m-1]})^{-1} \right|^{(1-kd)/2}} \cdot \frac{\exp\left(-\frac{1}{2} \text{tr}\left(S_t^{-1}(\Omega_t^*) \Omega_{t+1}^{-1}\right)\right)}{\exp\left(-\frac{1}{2} \text{tr}\left(S_t^{-1}\left(\Omega_t^{[m-1]}\right) \Omega_{t+1}^{-1}\right)\right)}$$

where:

$$S_t^{-1}(\Omega_t^*) = k \left(G^{1/2'}\right)^{-1} (\Omega_t^*)^d \left(G^{1/2}\right)^{-1}$$

$$S_t^{-1}\left(\Omega_t^{[m-1]}\right) = k \left(G^{1/2'}\right)^{-1} \left(\Omega_t^{[m-1]}\right)^d \left(G^{1/2}\right)^{-1}$$

Here, $\Omega_t^{[m-1]}$ denotes the current state of the Markov chain, while Ω_t^* is the new proposed value.

The last element of the history, Ω_T^{-1} , can be drawn directly from the Wishart distribution:

$$Wish\left(k+1, (S_{T-1}^{-1} + u_T u_T')^{-1}\right)$$

Step 4 - Drawing hyperparameters: G^{-1}

A random walk proposal is used:

$$(G^{-1})^* = (G^{-1})^{[m-1]} + W \quad (10)$$

where:

$$W = \begin{pmatrix} w_{11} & \cdots & w_{1p} \\ \vdots & \ddots & \vdots \\ w_{p1} & \cdots & w_{pp} \end{pmatrix}$$

$$\begin{pmatrix} w_{11} & w_{12} & \cdots & w_{pp} \end{pmatrix}' \sim N(0, \sigma_{G^{-1}}^2 I)$$

The matrix G^{-1} is assumed to follow a symmetric matrix variate normal distribution, and each of its $p(p+1)/2$ different elements can be drawn from a $p(p+1)/2$ -dimensional normal distribution.

The acceptance ratio can be written as:

$$AR = \frac{|G^*|^{-(k_0+kT-p-1)/2} \cdot \exp\left(-\frac{1}{2}tr\left(S_0^{-1}(G^*)^{-1}\right)\right)}{|G^{[m-1]}|^{-(k_0+kT-p-1)/2} \cdot \exp\left(-\frac{1}{2}tr\left(S_0^{-1}(G^{[m-1]})^{-1}\right)\right)} \cdot \frac{\exp\left(-\frac{1}{2}tr\left(\sum_{t=1}^T S_{t-1}^{-1}(G^*)\Omega_t^{-1}\right)\right)}{\exp\left(-\frac{1}{2}tr\left(\sum_{t=1}^T S_{t-1}^{-1}(G^{[m-1]})\Omega_t^{-1}\right)\right)}$$

where:

$$S_{t-1}^{-1}(G^*) = k \left((G^*)^{1/2'}\right)^{-1} \Omega_{t-1}^d \left((G^*)^{1/2}\right)^{-1}$$

$$S_{t-1}^{-1}(G^{[m-1]}) = k \left((G^*)^{1/2'}\right)^{-1} \Omega_{t-1}^d \left((G^*)^{1/2}\right)^{-1}$$

Step 5 - Drawing hyperparameters: k

Again, a random walk proposal is used for k :

$$k^* = k^{[m-1]} + \epsilon_k \quad (11)$$

$$\epsilon_k \sim N(0, \sigma_k^2)$$

and the acceptance ratio can be written as:

$$AR = \left(\frac{|k^* G^{-1}|^{-k^*/2}}{2^{pk^*/2} \prod_{i=1}^p \Gamma\left(\frac{k^*+1-i}{2}\right)} \right)^T \cdot \left(\frac{|k^{[m-1]} G^{-1}|^{-k^{[m-1]}/2}}{2^{pk^{[m-1]}/2} \prod_{i=1}^p \Gamma\left(\frac{k^{[m-1]}+1-i}{2}\right)} \right)^{-T} \cdot \frac{\prod_{t=1}^T |\Omega_{t-1}^{-1}|^{-dk^*/2} |\Omega_t^{-1}|^{k^*/2}}{\prod_{t=1}^T |\Omega_{t-1}^{-1}|^{-dk^{[m-1]}/2} |\Omega_t^{-1}|^{k^{[m-1]}/2}} \cdot \frac{\prod_{t=1}^T \exp\left(-\frac{1}{2} \text{tr}\left(S_{t-1}^{-1}(k^*) \Omega_t^{-1}\right)\right)}{\prod_{t=1}^T \exp\left(-\frac{1}{2} \text{tr}\left(S_{t-1}^{-1}(k^{[m-1]}) \Omega_t^{-1}\right)\right)}$$

with

$$S_{t-1}^{-1}(v^*) = v^* \left(G^{1/2'}\right)^{-1} \Omega_{t-1}^d \left(G^{1/2}\right)^{-1}$$

$$S_{t-1}^{-1}(v^*) = v^* \left(G^{1/2'}\right)^{-1} \Omega_{t-1}^d \left(G^{1/2}\right)^{-1}$$

Step 6 - Drawing hyperparameters: d

Finally, a random walk proposal is adopted for d as well:

$$d^* = d^{[m-1]} + \epsilon_d \quad (12)$$

$$\epsilon_d \sim N(0, \sigma_d^2)$$

the acceptance ratio is:

$$AR = \frac{\prod_{t=1}^T |\Omega_{t-1}|^{kd^*/2}}{\prod_{t=1}^T |\Omega_{t-1}|^{kd^{[m-1]}/2}} \cdot \frac{\prod_{t=1}^T \exp\left(-\frac{1}{2} \text{tr}\left(S_{t-1}^{-1}(d^*) \Omega_t^{-1}\right)\right)}{\prod_{t=1}^T \exp\left(-\frac{1}{2} \text{tr}\left(S_{t-1}^{-1}(d^{[m-1]}) \Omega_t^{-1}\right)\right)}$$

and

$$S_{t-1}^{-1}(d^*) = v \left(G^{1/2'} \right)^{-1} \Omega_{t-1}^{d^*} \left(G^{1/2} \right)^{-1}$$

$$S_{t-1}^{-1} \left(d^{[m-1]} \right) = v \left(G^{1/2'} \right)^{-1} \Omega_{t-1}^{d^{[m-1]}} \left(G^{1/2} \right)^{-1}$$

In practice, I found that the algorithm is much more efficient in terms of convergence if J steps of the *MH* procedure are performed for each draw of B^T and Q (with J being a small number, for instance 3 or 5). Thus, the sampler can be summarized as follows:

1. Initialize Ω^T and Q .
2. Sample B^T from $p(B^T | y^T, \Omega^T, Q, G, k, d) = p(B^T | y^T, \Omega^T, Q)$.
3. Sample Q from $p(Q | y^T, B^T, \Omega^T, G, k, d) = p(Q | y^T, B^T)$.
4. Sample Ω_T^{-1} from $Wish(k+1, (S_{T-1}^{-1} + u_T u_T')^{-1})$. Then, sample each proposed value $(\Omega_t^*)^{-1}$ from $Wish(k, \tilde{S}_{t-1})$. If the proposed value is not accepted, set the draw of Ω_t equal to $\Omega_t^{[m-1]}$.
5. Sample the proposed value G^* from the random walk distribution (10); if G^* is not accepted, set the draw of G equal to $G^{[m-1]}$, i.e. the previous value of the algorithm.
6. Sample the proposed value k^* from the random walk distribution (11); if k^* is not accepted, set the draw of k equal to $k^{[m-1]}$.
7. Sample the proposed value d^* from the random walk distribution (12); if d^* is not accepted, set the draw of d equal to $d^{[m-1]}$.
8. Repeat 4 – 7 for J times;
9. Go back to #2

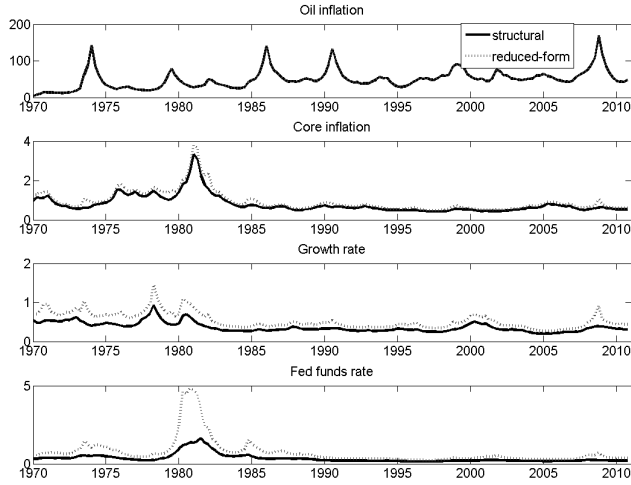
References

- [1] Baumeister, C., Peersman, G., 2008. Time-Varying Effects of Oil Supply Shocks on the US Economy. Mimeo, Faculty of Economics and Business Administration, Ghent University.
- [2] Bernanke, B.S., Gertler, M., Watson, M., 1997. Systematic monetary policy and the effects of oil price shocks. *Brookings Papers on Economic Activity* (1), 91–142.
- [3] Bernanke, B.S., Gertler, M., Watson, M., 2004. Oil shocks and aggregate macroeconomic behavior: the role of monetary policy: reply. *Journal of Money, Credit, and Banking*, vol. 36, no. 2, 287-91.
- [4] Blanchard, O. J., Gali, J., 2007. The Macroeconomic Effects of Oil Shocks: Why are the 2000s So Different from the 1970s? National Bureau of Economic Research, Working Papers: 13368.
- [5] Canova, F., Gambetti, L., 2009. Structural changes in the US economy: Is there a role for monetary policy? *Journal of Economic Dynamics and Control*, vol. 33(2), pages 477-490.
- [6] Carter, C.K., Kohn, R., 1994. On Gibbs Sampling for State Space Models. *Biometrika*, 81, 541-553.
- [7] Chib, S., Greenberg, E., 1995. Understanding the Metropolis-Hasting algorithm. *The American Statistician*, vol. 49, n.4.
- [8] Clark, T.E., Terry, S.J., 2010. Time Variation in the Inflation Passthrough of Energy Prices, *Journal of Money, Credit and Banking*, vol. 42, no. 7, 1419–1433.
- [9] Cogley, T., Sargent, T.J., 2001. Evolving post world war II US inflation dynamics. In: *NBER Macroeconomics Annual*, vol. 16, pp. 331–373.
- [10] Cogley, T., Sargent, T.J. 2005a. Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S. *Review of Economic Dynamics* 8, 262-302.
- [11] Hamilton, J. D., Herrera A. M., 2004. Oil shocks and aggregate macroeconomic behavior: the role of monetary policy: comment. *Journal of Money, Credit, and Banking*, vol. 36, no. 2, April, pp. 265-86.

- [12] Hasting, W.K., 1970. Monte Carlo Sampling Methods Using Markov Chains and Their Applications. *Biometrika*, 57, 97-109.
- [13] Herrera, A., Pesavento, E., 2009. Oil Price Shocks, Systematic Monetary Policy, And The "Great Moderation". *Macroeconomic Dynamics*, vol. 13, n. 1, February, pp. 107-137.
- [14] Kilian, L., 2008. The Economic Effects of Energy Price Shocks. *Journal of Economic Literature*, vol. 46, n.4, December, pp. 871-909.
- [15] Kilian, L., Lewis, L., 2009. Does the Fed Respond to Oil Price Shocks?. C.E.P.R. Discussion Papers: 7594.
- [16] Kilian, L., Murphy, D., 2010. Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models. MPRA Paper 80210, University Library of Munich, Germany.
- [17] Lippi, F., Nobili, A., 2008. Oil and the Macroeconomy: A Structural VAR Analysis with Sign Restrictions. CEPR Discussion Papers n.6830.
- [18] Philipov, A., Glickman, M., 2006. Multivariate Stochastic Volatility via Wishart Processes. *Journal of Business and Economic Statistics*, 24 (3), 313-328.
- [19] Primiceri, G., 2005. Time Varying Structural Vector Autoregression and Monetary Policy. *Review of Economic Studies* 72, 821-852.
- [20] Rinnergschwentner, W., Tappeiner, G., Walde, J., 2011. Multivariate Stochastic Volatility via Wishart Processes - A Continuation. Working Papers 2011-19, Faculty of Economics and Statistics, University of Innsbruck.
- [21] Rotemberg, J. J., Woodford, M., 1996. Imperfect competition and the effects of energy price increases. *Journal of Money, Credit, and Banking* 28, 549-577.
- [22] Rubio-Ramirez, J., Waggoner, D., Zha, T., 2010. Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference. *Review of Economic Studies*, vol. 77(2), pp. 665-696.


Figures

Figure 1 - Standard deviations, reduced-form and structural model



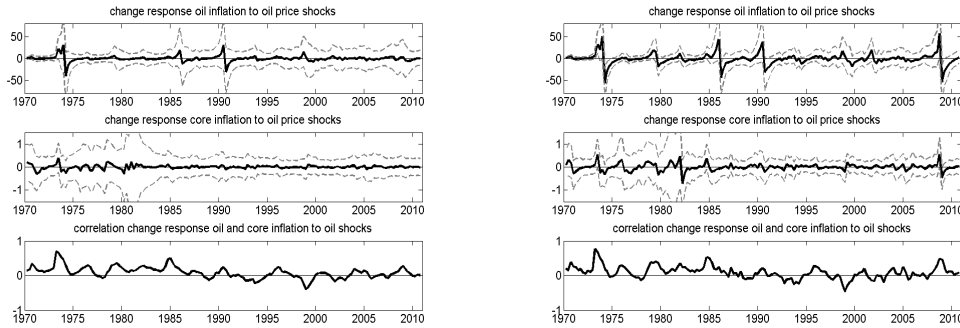
Note: Median value of the standard deviations of the reduced-form innovations (dotted line) and of the shocks in the structural VAR model identified using ordering restrictions (continuous line). These values were computed from the draws of Ω^T obtained from the MCMC procedure described in the text.

Figure 2 - Response of core inflation to oil price shocks



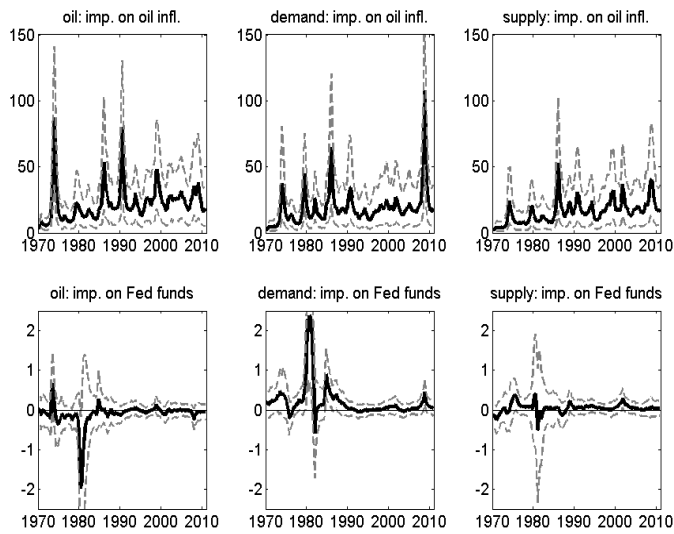
Note: Median response of core inflation to oil price shocks, together with the 16th and 84th percentile bands (panel 1) and number of histories A^T in which this response is positive (panel 2).

Figure 3 - Change response of oil inflation and core inflation to oil price shocks



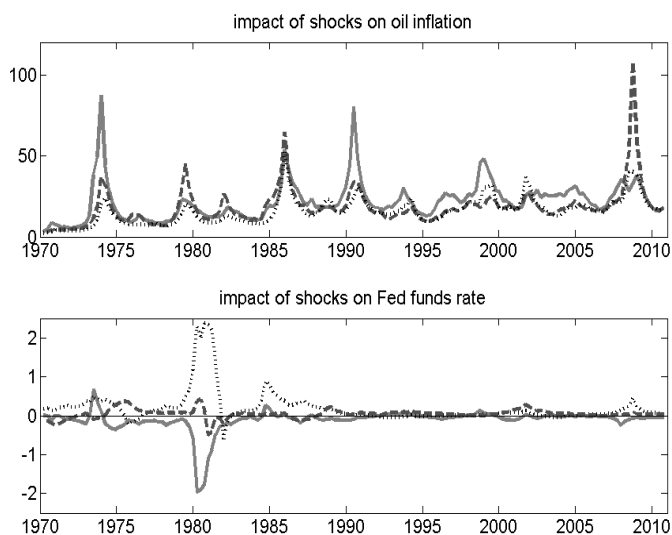
Note: Median response of oil inflation and core inflation to oil price shocks, together with the 16th and 84th percentile bands, and correlations between these responses. The left panel reports the results for the model identified using sign restrictions, while the right panel reports the results for the model identified using the Cholesky decomposition.

Figure 4 - Response of oil inflation and Fed funds to shocks



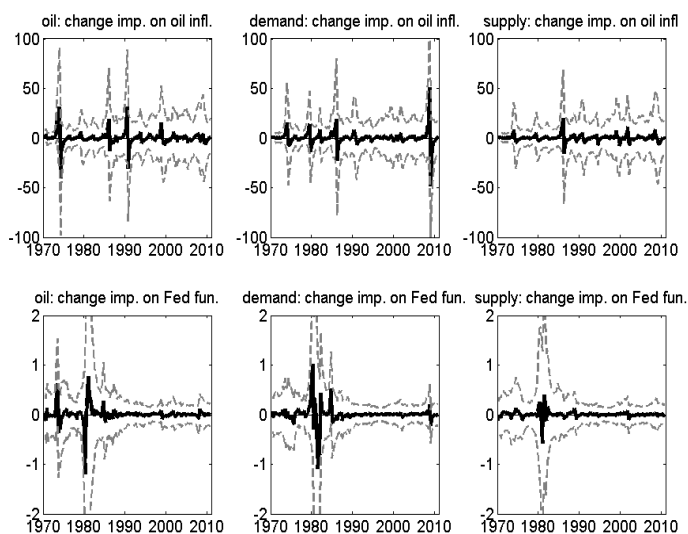
Note: The top panel reports the median response of oil inflation to oil price shocks, US demand shocks, and US supply shocks, together with the 16th and 84th percentile bands. The bottom panel reports the median response of the Fed funds to oil price shocks, US demand shocks, and US supply shocks, together with the 16th and 84th percentile bands.

Figure 5 - Response of oil inflation and Fed funds to shocks




Note: The top panel reports the median response of oil inflation to oil price shocks (continuous line), US demand shocks (dashed line), and US supply shocks (dotted line). The bottom panel reports the median response of the Fed funds to oil price shocks (continuous line), US demand shocks (dashed line), and US supply shocks (dotted line). For clarity of exposition, the 16th and 84th percentile bands are not reported.

Figure 6 - Changes in the response of oil inflation and Fed funds to shocks



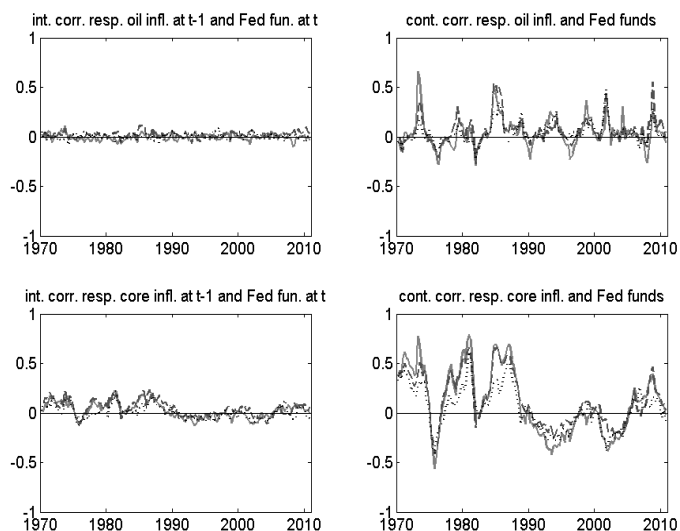
Note: The top panel reports the median change in the response of oil inflation to oil price shocks, US demand shocks, and US supply shocks, together with the 16th and 84th percentile bands. The bottom panel reports the median change in the response of the Fed funds to oil price shocks, US demand shocks, and US supply shocks, together with the 16th and 84th percentile bands.

Figure 7 - Changes in the response of oil inflation and Fed funds to shocks



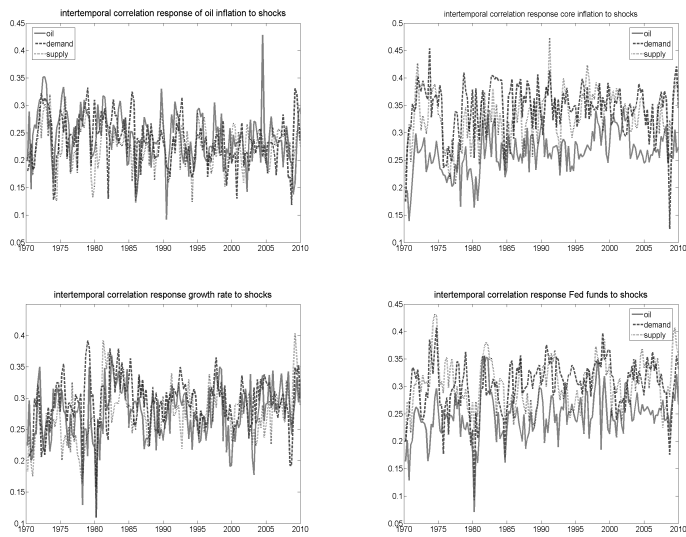
Note: The top panel reports the median change in the response of oil inflation to oil price shocks (continuous line), US demand shocks (dashed line), and US supply shocks (dotted line). The bottom panel reports the median change in the response of the Fed funds to oil price shocks (continuous line), US demand shocks (dashed line), and US supply shocks (dotted line). For clarity of exposition, the 16th and 84th percentile bands are not reported.

Figure 8 - Intertemporal correlations selected responses to oil price shocks



Note: The left-hand panel reports the intertemporal correlations between the impact of oil price shocks (continuous line), US demand shocks (dashed line), and US supply shocks (dotted line) on oil inflation and core inflation at time $t - 1$ and the Federal funds rate at time t . The right-hand panel reports the contemporaneous correlations between the same variables.

Figure 9 - Intertemporal correlations responses to shocks



Note: Intertemporal correlations in the response of the endogenous variables (oil inflation, core inflation, growth rate, and Federal funds rate) to the different shocks (oil, continuous line; US demand, dashed line; US supply, dotted line).

Tables

Table 1 - Distribution of selected parameters in the retained draws

	16 th perc.	median	84 th perc.	accept. rate
g_{11}	2.881	3.340	3.865	0.368
g_{22}	0.508	0.574	0.655	0.368
g_{33}	0.425	0.499	0.574	0.368
g_{44}	0.446	0.511	0.587	0.368
k	6.766	7.709	9.108	0.330
d	0.771	0.810	0.845	0.467

Table 2 - Median values of G^{-1}

	oil infl.	core infl.	growth rate	Fed funds
oil infl.	3.340	0.475	0.251	-0.001
core infl.	0.475	0.574	0.004	-0.007
growth rate	0.251	0.004	0.499	0.010
Fed funds	-0.001	-0.007	0.010	0.511