Heteroskedasticity as a Complementary Identification Strategy

Andrea Carriero, Massimiliano Marcellino and Tommaso Tornese *

VERY PRELIMINARY (June 2021)

Abstract

We show how to combine heteroskedasticity with sign restrictions and external instruments in order to sharpen identification in structural VARs.

We overcome the limitations of the identification strategies based on sign restrictions or external instruments, i.e. potentially large identified sets, taking advantage of the information introduced by heteroskedasticity. Moreover, we overcome the labeling problem inherent in identification through heteroskedasticity, considering additional identifying information in the form of sign restrictions or external instruments.

We illustrate the implementation of the procedure for the identification of various shocks and their effects. Specifically, we consider labour market, oil market, monetary and fiscal policy shocks.

Key words: SVAR, Identification, Heteroskedasticity, Sign restrictions, Proxy variables.
Journal-classification: C11, C32, D81, E32.

*Andrea Carriero: Queen Mary University of London and University of Bologna; e-mail: a.carriero@qmul.ac.uk. Massimiliano Marcellino: Bocconi University; e-mail: massimiliano.marcellino@unibocconi.it. Tommaso Tornese: Queen Mary, University of London; e-mail: t.tornese@qmul.ac.uk

1 Introduction

Since the seminal contribution of (Sims, 1980), the empirical analysis of shocks propagation is most often carried out using Vector Autoregressive (VAR) models. VARs can offer a very good description of the dynamics of macroeconomic and financial data, and - when coupled with an appropriate prior - have shown a remarkably good performance in out of sample forecasting. If one wants to use VARs to answer policy questions, to analyze the effects of policy shocks, or more generally to give a causal interpretation to such class of models, it becomes necessary to map the VAR into a system of simultaneous equations.

As is well known in the mcroeconometrics literature, the main difficulty of this task comes from the fact that such a mapping is not unique, and one has to rely on assumptions derived from economic theory or on external information to restrict the set of admissible structural models that share the same reduced form representation. The fundamental problem is that the data are only informative about the joint dynamics of macroeconomic and financial variables, but they cannot tell anything about the causal direction in which the underlying structural relations between economic variables manifest. In other words, VAR models are silent about the contemporaneous causal links in economic variables unless further assumptions are made. Traditionally, these additional assumptions take the form of zero restrictions on the matrix that defines the contemporaneous relation across the variables, or the matrix measuring the effect of a structural shock on a certain variable, either on impact or in the long run.

However, the assumption that a shock does not affect at all certain macroeconomic variable is too restrictive in many cases. Moreover, the number of zero restrictions required increases rapidly with the size of the model, making restrictions even harder to choose and to defend. For this reason researchers have endeavored to devise alternative identification strategies able to narrow down the set of structural models in which a VAR can be mapped (see Kilian and Lutkepohl (2017) for a thorough treatment). Two particularly successful strategies are sign restrictions and the use of external instruments (proxies).

The sign restriction approach, formulated by Faust (1998), Canova and De Nicolò (2002), and Uligh (2005), is based on the conjecture that the effect of a shock on a given variable, at some horizon, has a pre-specified sign. More recently, this idea has been revisited by Baumeister and Hamilton (2015), who impose sign restrictions on the structural contemporaneous relationships between variables on the basis of widely accepted theoretic postulates. The remarkable merit of this class of identification strategies is that it replaces the heroic assumptions on which zero restrictions are based with arguably milder conjectures about the sign of macroeconomic relations. The price for being able to impose these milder conditions, however, is potentially high: Sign restrictions provide set identification, as opposed to point identification. This means that, contrary to what is usually achieved imposing zeros restrictions, the resulting set of admissible structural models associated with the reduced form VAR is not singleton, and inference must be drawn on the basis of the whole set.

The external instruments approach, pioneered by Stock (2008), Stock and Watson (2012) and Mertens and Ravn (2013), on the other hand, exploits the availability of variables that are believed to be correlated with some structural shocks of interest, but uncorrelated with all the other shocks. Also this approach allows to avoid the imposition of eccessively strong restrictions on the model and is being widely and succesfully implemented. However, also this strategy can (and does) in some cases run into problems. In particular, when more than one proxy variable is available for more than one shock, also this approach can only yield set identification, and some researchers in some instances had to complement this identification strategy with the reintroduction of some zero restrictions (e.g. Mertens and Ravn (2013), Lakdawala (2019)).

Parallel to these two approaches (which are used widely), a separate strand of the recent literature has pointed towards exploiting heteroskedasticity for identification. Notable examples are Sentana and Fiorentini (2001), Rigobon (2003), Rigobon and Sack (2004) and Lanne et al. (2010), who show that, if changes in the variances of structural shocks are heterogeneous enough, reduced form information is sufficient to pin down the parameters of the structural model. Lewis (2018a, 2018b) extends these ideas to more general and unknown variance processes, showing in a non-parametric setup that the conditions required for identification are generally mild, and that identification can be achieved even in presence of misspecification of the variance process. The important limitation that has prevented identification through heteroskedasticity from being widely used by empirical researchers is the absence of labeling of the identified shocks. In fact, the presence of heteroskedasticity allows the econometrician to determine uniquely the structural shocks that drive economic fluctuations, but it does not convey information about what is the nature of the identified innovations; one can only attach a label to shocks based on the effects it produces on the variables of the system, as done, for instance, by Brunnermeier et al. (2021).

In this paper we argue that the identifying information embedded in heteroskedasticity can be used to sharpen the identification schemes based on sign restrictions and external instruments. We show how heteroskedasticity can be routinely and effectively used to reduce the identified set of VARs identified via sign-restrictions and external-instruments approaches. Similarly, we show that the information contained in sign-restrictions and external-insturments schemes can help overcoming the key issue inherent in identification through heteroskedasticity, i.e. the labeling problem. In summary, we support a view that researchers should consider blending together these alternative identification strategies as the blend contains more idenfifying information than the parts.

We illustrate the convenience of this blended approach by revisiting empirical results presented in seminal papers based on either sign restrictions or proxy variables. In particular, we reconsider the applications in Baumeister and Hamilton (2015, 2018, 2019), analyzing the effects of labour market, oil market and monetary policy shocks respectively. Moreover, we reproduce the Proxy-SVAR model of Mertens and Ravn (2013), whose aim is to assess the effects of APITR (Average Personal Income Tax Rate) and ACITR (Average Corporate Income Tax

Rate) shocks, and we do this by relaxing the zero restrictions that allow them to achieve point identification.

Our empirical applications show that the introduction of identifying information coming from heteroskedasticity can generate substantial changes in the results. For example, in the labor market application, it leads to a significantly steeper aggregate labour demand and a supply that is almost flat. Related to this, the responses of wages and employment to demand shocks are dramatically revised, they are both less responsive when compared to the case with constant variances. In the application on oil, the results with heteroskedasticity support Kilian and Murphy's (2012) assumption that the price elasticity of supply is smaller than 0.0258, which has been harshly criticized by Baumeister and Hamilton with the homoskedastic model. Moreover, the world demand for oil becomes much steeper and inventories have an insignificant sensitivity to the quantity of crude oil produced, while in the standard SVAR this sensitivity is negative and significant. In the monetary SVAR, a monetary contraction has more recessive and persistent consequences with heteroskedasticity, and there is a significant drop of hourly earnings after two years from the shock. Other shocks (e.g., demand and supply) also have different effects when allowing for heteroskedasticity in the SVAR. In the fiscal policy application, both types of fiscal contractions have deep negative effects on output once heteroskedasticity is taken into consideration. Moreover, following a corporate income tax shock, the corporate income tax base actually decreases. This mechanism, together with the negative effect on output, leads government debt to rise in response to the unexpected corporate tax increase. On the other hand, a personal income tax shock raises the tax base and allows a reduction of the government debt, as one would expect.

The paper is structured as follows: Section 2 describes the structural VAR (SVAR) model and discusses the identification problem. Section 3 presents the class of sign restricted SVAR introduced by Baumeister and Hamilton (2015), illustrates the Bayesian algorithm we employ in order to draw from the posterior of a heteroskedastic version of their model (SVAR-H) and revisit the empirical results the authors obtain applying their methodology in a series of works. Section 4 describes the Proxy-SVAR approach, adapts it to account for heteroskedasticity and re-examines the effects of personal and corporate tax surprises using data from Mertens and Ravn (2013). Section 4 summarizes and concludes.

2 Empirical Strategy

2.1 SVAR Model and Identification

The joint dynamics of macroeconomic variables is often modeled as a VAR(p) process:

$$y_t = \Pi x_{t-1} + u_t, \tag{1}$$

where y_t is a $n \times 1$ vector of endogenous variables, x_{t-1} is the $k \times 1$ ($k = np + 1$) vector

$\left[1, y'_{t-1}, \ldots, y'_{t-p}\right]'$, and Π is a $n \times k$ matrix of coefficients. The $n \times 1$ vector of prediction errors u_t is assumed to be normally distributed with zero mean and variance Ω, and it is believed to be a linear combination of structural shocks: $u_t = A^{-1}\varepsilon_t$, $\varepsilon_t \sim N(0, \Lambda)$, with $\Lambda = diag[\lambda_1, ..., \lambda_n]$ and A (or A^{-1}) potentially full with ones on the main diagonal.

Most often, researchers are interested in the propagation of structural innovations, which are considered the drivers of economic fluctuations. For this reason, policy analysis is usually carried out considering the structural form of the VAR model. This is obtained premultiplying both sides of equation (1) by A:

$$Ay_t = Bx_{t-1} + \varepsilon_t, \tag{2}$$

where $B = A\Pi$.

The main difficulty when moving from the reduced form VAR in (1) to the structural representation in (2) is that this mapping is not unique, unless more information is introduced in the model. This problem is known as the identification problem in the SVAR literature and it arises because the data are only informative about the parameters (Π, Ω), but they cannot distinguish between alternative structural specifications, (A, B, Λ), that give rise to the same reduced form model. To see this, note that the following must hold:

$$\Omega = A^{-1}\Lambda\left(A^{-1}\right)' \tag{3}$$

Since the reduced from variance Ω is symmetric, (3) provides a system of $n(n+1)/2$ equations in which the left hand side can be inferred from the data, but the right hand side contains $n(n-1)+n$ unknown elements. As a result, (3) does not have a unique solution and more information has to be introduced in order to obtain an economically meaningful parametrization of the SVAR (2).

This paper focuses on three main ways of introducing such identifying information in the model, namely: sign restrictions, external instruments and heteroskedasticity.

The use of sign restrictions consists in considering only the solutions of equation (3) that satisfy the assumption that certain elements of A^{-1}(or A in the case of Baumaister and Hamilton, 2015, 2018, 2019) have a prespecified sign. These assumptions are usually made on the basis of widely acceptable theoretic justifications, which explains why this identification strategy has become one of the most popular in the empirical literature. Such mild assumptions, however, are often not enough to make the solution of (3) unique. For this reason, inference must be drawn from a set of equally likely parametrizations (A, B, Λ) with possibly vague interpretations.

External instruments, when available, provide precious information that can be used to obtain identification. The special feature of proxy variables is that they are known to be correlated with the structural shocks of interest (*relevance condition*) but uncorrelated with all other shocks (*exogeneity condition*).

If one proxy is available for the i-th shock, the parameters of the i-th equation in the SVAR

(2) can be uniquely determined. However, in practice it is often the case that $m > 1$ instruments are correlated with m shocks of interest and uncorrelated with all the other $n-m$. In that case, the parameters of the m equations in (2) that are associated with the m innovation of interest are only set identified, in the sense that there exist a set of parametrizations that satisfy the conditions implied by the external instruments, but all the elements of this set are observationally equivalent. In order to reduce the set to a collection of economically meaningful models, researchers often impose additional zero restrictions, as in Mertens and Ravn (2013) and Lakdawala (2019), or sign restrictions, as in Piffer and Podstawski (2017) and Braun and Bruggemann (2020).

Identification through heteroskedasticity, on the other hand, is a purely statistical method to achieve a unique parametrization (A, B, Λ). The strategy exploits the presence of time variation in the variances of structural innovations ε_t, which translates in time variation of the variance/covariance of the reduced form errors u_t.

We will assume for the rest of the paper that these variances evolve as a regime switching process, where the dates of regime changes are known to the researcher. In addition to this being often the case in practice, knowing the exact date of regime changes is not necessary for consistency. Rigobon (2003) and Sims (2020) show that estimates of the structural parameters are still consistent if the windows of heteroskedasticity are misspecified.

Intuitively, heteroskedasticity adds more equations than free parameters to the system in (3). If one considers the case with S regimes, the mapping of the reduced form variance Ω_s, $s = 1, \ldots, S$, into the structural parameters (A, Λ_s) becomes:

$$\Omega_s = A^{-1}\Lambda_s\left(A^{-1}\right)' \tag{4}$$

which is now a system of $S[n(n+1)/2]$ equations in $n(n-1)+Sn$ unknowns. Results in Lanne et al. (2010) and Lutkepohl and Wozniak (2020) prove that, as long as there exist two regimes $s_1, s_2 \in \{1, \ldots, S\}$ for which $\lambda_{i,s_1}/\lambda_{j,s_1} \neq \lambda_{i,s_2}/\lambda_{j,s_2}$, $\forall j \neq i$ and $\forall i \in \{1, \ldots, n\}$, then the parameters of the structural model (2) are uniquely identified up to a permutation of the columns of A^{-1}. In other words, if changes in the variance of shocks are not proportional across all regimes, heteroskedasticity allows to achieve point identification but it is silent about the economic nature of ε_t, because any ordering of the shocks is equally likely.

The main assumption behind identification through heteroskedasticity is that the impact of shocks on economic variables, defined by A^{-1}, is time invariant. Although this assumption is not completely innocuous, it is also present in the other identification strategies discussed above, therefore it cannot be considered a discriminator between the use of sign restrictions, proxy variables or heteroskedasticity for identification.

In this paper we argue that there is a definite gain to be obtained by introducing heteroskedaticity in SVAR identified via sign restrictions or external instruments. There is an obvious "reduced form" gain which will come from the fact that a model with time varying volatility is in most cases a better representation of the data, which reduces the risk of model misspecification and generally improves the efficiency of the estimators. And there is also a "structural form" gain insofar modelling the variances as time varying allows to extract more identifying information from the data, sharpening the identification of the structural parameters.

3 Sign Restrictions and Heteroskedasticity

The traditional way of implementing sign restrictions is to estimate reduced form parameters and then to consider a uniform prior over the space of orthogonal matrices Q that map the reduced form model into a structural model that satisfies the desired sign restrictions on A^{-1}. The idea is to explore the set of all the structural parametrizations compatible with the estimated (Π, Ω) and with sign restrictions, by drawing rotations of the innovations ε_t associated with different matrices Q.[1] Note, in fact, that (3) continues to hold if we rotate the structural shocks using an orthogonal matrix ($QQ' = I$):

$$\Omega = A^{-1}\Lambda^{1/2}QQ'\left(\Lambda^{1/2}\right)'\left(A^{-1}\right)' \tag{5}$$

In an important paper, Baumeister and Hamilton (2015) point out some relevant limitations of this approach. Their main argument is that the prior over the admissible SVAR parametrizations is never updated by the data, and apparently uninformative priors for the rotation matrix Q can sometimes translate into very informative and implausible priors for objects of interest, such as Impulse-Response Functions (IRF).

Their proposed method to overcome these limitations follows from the idea of Sims and Zha (1998) of estimating the structural form in (2) directly. Baumeister and Hamilton (2015) suggest to introduce sign restrictions in this setting by placing truncated stundent-t priors on specific elements of A.

Although it is not an essential aspect of their suggested method, Baumeister and Hamilton's preference for considering sign restrictions on the contemporaneous relations embedded in A, rather than on the impact matrix A^{-1}, comes from the fact that the former type of constraints are more easily elicited on the basis of economic theory.

For example, in order to identify the simultaneous equations describing demand and supply in SVAR models, the researcher can impose that price and quantity are negatively related in the former and positively related related in the latter.

Furthermore, it must be noted that, contrary to traditional sign restricted SVARs, where the structural parameters are only set identified, in the class of models proposed by Baumeister and Hamilton (2015), the posterior distribution of the structural parameters provides point identification, because the informative priors introduced on the elements of A revise the shape of the flat regions of the likelihood.

[1] Rubio-Ramirez, Waggoner and Zha (2010) develop an efficient algorithm to perform this task.

In this section, we adopt Baumeister and Hamilton's approach and show how heteroskedasticity can be introduced in the SVAR to make the data more informative about the structural parameters. In this way, as long as the conditions stated in Lanne et al. (2010) and Luktepohl and Wozniak (2020) are satisfied, there will be no regions of the parameters space in which the likelihood is flat.

Subsection 3.1 describes the general Bayesian algorithm we use to draw from the posterior distribution of the SVAR-H model, while in subsections 3.2, 3.3 and 3.4 we revisit the empirical results in Baumaister and Hamilton (2015, 2018, 2019) through the lenses of our heteroskedastic model.

3.1 SVAR-H

In this section we borrow notation from Luktepohl and Wozniak (2020) and write $\lambda_{i,s} = \lambda_{i,1}\omega_{i,s}$, $s = 1,...,S$, where it is understood that $\omega_{i,1} = 1$ for every $i \in \{1,...,n\}$. Although this specification does not change any property of the model, it provides an immediate way to check that enough heterogeneity is present in the changes of variance regimes.

If we assume that the data are generated by a SVAR(p) model with heteroskedastic structural innovations, $\varepsilon_t \sim N(0,\Lambda_s)$, $s = 1,...,S$, the likelihood is then:

$$p(y_{1:T}|A,B,\Lambda_{1:S}) = (2\pi)^{-Tn/2}|\det(A)|^T \left(\prod_{i=1}^{n}\lambda_{i,1}^{-T/2}\right)\left(\prod_{i=1}^{n}\prod_{s=1}^{S}\omega_{i,s}^{-T_s/2}\right)\times$$
$$\exp\left\{-\frac{1}{2}\sum_{i=1}^{n}\lambda_{i,1}^{-1}\left[\sum_{s=1}^{S}\omega_{i,s}^{-1}\left(A_i y^{(s)} - B_i x^{(\mathbf{s})}\right)\left(A_i y^{(s)} - B_i x^{(\mathbf{s})}\right)'\right]\right\} \tag{6}$$

where T is the time series size of the sample and T_s is the length of regime s. The matrices $y^{(s)}$ and $x^{(s)}$ collect the values of y_t and x_{t-1} observed during regime s and have size $n\times T_s$ and $k\times T_s$ respectively, while A_i and B_i denote the i-th row of A and B.

If we assume student-t priors for the n_α free elements in A, α, conditionally normal priors with mean μ_i and variance V_i on the vectors of coefficients B_i, independent inverse-gamma priors with $d_{i,1}/2$ degrees of freedom and scale parameter $\zeta_{i,1}/2$ for the variances $\lambda_{i,1}$, inverse-gamma priors with $d_{i,s}/2$ degrees of freedom and scale parameter $\zeta_{i,s}/2$ for the ratios $\omega_{i,s}$, for $s = 2,...,S$, it is shown in the appendix that the conditional posterior distribution of B_i is normal with moments:

$$\bar{V}_i = \left[V_i^{-1} + \lambda_{i,1}^{-1}\sum_{s=1}^{S}\omega_{i,s}^{-1}x^{(s)}\left(x^{(s)}\right)'\right]^{-1}$$
$$\bar{\mu}_i' = \left[\mu_i' V_i^{-1} + \lambda_{i,1}^{-1}\sum_{s=1}^{S}\omega_{i,s}^{-1}A_i y^{(s)}\left(x^{(s)}\right)'\right]\bar{V}_i \tag{7}$$

The variances $\lambda_{i,1}$, conditional on all other parameters, have inverse-gamma posteriors with degrees of freedom and scale parameters:

$$\frac{\bar{d}_{i,1}}{2} = \frac{d_{i,1} + T}{2}$$

$$\frac{\bar{\zeta}_{i,1}}{2} = \frac{\zeta_{i,1} + \sum_{s=1}^{S} \omega_{i,s}^{-1} \left(A_i y^{(s)} - B_i x^{(\mathbf{s})}\right) \left(A_i y^{(s)} - B_i x^{(\mathbf{s})}\right)'}{2} \tag{8}$$

The ratios $\omega_{i,s}$ have inverse-gamma conditional posteriors with degrees of freedom and scale parameters:

$$\frac{\bar{d}_{i,s}}{2} = \frac{d_{i,s} + T_s}{2}$$

$$\frac{\bar{\zeta}_{i,s}}{2} = \frac{\zeta_{i,s} + \lambda_{i,1}^{-1} \left(A_i y^{(s)} - B_i x^{(\mathbf{s})}\right) \left(A_i y^{(s)} - B_i x^{(\mathbf{s})}\right)'}{2} \tag{9}$$

Finally, the vector of free elements of A, α, has a non standard posterior distribution whose density is:

$$p\left(\alpha | Y, B, \Lambda_{1:S}\right) \propto p(\alpha) \left|\det(A)\right|^{T} \exp\left\{ -\frac{1}{2} \sum_{i=1}^{n} \lambda_{i,1}^{-1} \left[\sum_{s=1}^{S} \omega_{i,s}^{-1} \left(A_i y^{(s)} - B_i x^{(\mathbf{s})}\right) \left(A_i y^{(s)} - B_i x^{(\mathbf{s})}\right)' \right] \right\} \tag{10}$$

where $p(\alpha)$ is the product of the univariate student-t densities placed on the single free elements in A.

Given these conditional densities, draws from the joint posterior of the parameters of the SVAR-H model can be obtained implementing a Markov Chain Monte Carlo algorithm that cycles through the following steps for $m = 1, ..., M$, where M is a sufficiently large number of iterations[2]:

1. Draw from $p\left(\Lambda_s | Y, A^{(m-1)}, B^{(m-1)}\right)$ for $s = 1, ..., S$, which is accomplished by:
   (a) drawing from $p\left(\lambda_{i,1} | Y, A_i^{(m-1)}, B_i^{(m-1)}, \omega_{i,2:S}^{(m-1)}\right)$ for $i = 1, ..., n$;
   (b) drawing from $p\left(\omega_{i,s} | Y, A_i^{(m-1)}, B_i^{(m-1)}, \lambda_{i,1}^{(m)}\right)$ for $i = 1, ..., n$ and for $s = 2, ..., S$;
2. Draw from an approximation of $p\left(\alpha | Y, B^{(m-1)}, \Lambda_{1:S}^{(m)}\right)$ using a Random Walk Metropolis-Hastings algorithm, which:
   (a) draws a candidate $\alpha^{\star} = \alpha^{(m-1)} + \xi \nu$, where $\nu \sim t(2)$ and has dimension $n_\alpha \times 1$, while ξ is a tuning parameter that guarantees an acceptance rate between 30% and

[2] In our empirical applications $M = 2,000,000$ with a burn-in of $M/2$

40%[3];

(b) compute $\vartheta = \min\left\{ \frac{p\left(\alpha^{\star}|Y, B^{(m-1)}, \Lambda_{1:S}^{(m)}\right)}{p\left(\alpha^{(m-1)}|Y, B^{(m-1)}, \Lambda_{1:S}^{(m)}\right)}, 1 \right\}$, and set $\alpha^{(m)} = \alpha^{\star}$ with probability ϑ and $\alpha^{(m)} = \alpha^{(m-1)}$ with probability $(1-\vartheta)$;

3. Draw from $p\left(B_i|Y, A_i^{(m)}, \lambda_{i,1:S}^{(m)}\right)$, for $i = 1, ..., n$.

The main difference between our algorithm and the one implemented by Baumeister and Hamilton (2015, 2018, 2019) is that the presence of heteroskedasticity does not allow us to have natural conjugate priors, and, hence, forces us to resort to MCMC methods in the place of the plain Monte Carlo sampling used in their applications. As a result, the prior distributions we consider in the next subsections will be close but not identical to the one assigned in the respective original papers.

To maximize the reliability of our results, all the comparisons we make throughout the paper are based on the homoskedastic version of our specification, obtained setting $\omega_{i,s} = 1$ for all i and s. In this way, the differences between the estimated homoskedastic and heteroskedastic models can only be driven by information conveyed by the data, and cannot be affected by differences in parameters priors.

3.2 Labour Market Shocks

The first empirical example in which Baumeister and Hamilton (2015) illustrate their proposed method is a bivariate SVAR describing aggregate labour demand and supply in the U.S.

The model takes the form in (2) with:

$$A = \begin{bmatrix} -\beta^d & 1 \\ -\alpha^s & 1 \end{bmatrix} \tag{11}$$

and $y_t = [\Delta w_t, \Delta n_t]$, where w_t and n_t are 100 times the logarithm of the wage level and employment level, while Δ denotes the quarterly difference. [4]

The free elements of A, $\beta^d$and α^s, are interpreted as the short run demand and supply elasticity. The authors place on β^d a student-t prior with 3 degrees of freedom, location parameter $c_\beta = -0.6$ and scale parameter $\sigma_\beta = 0.6$, and they set for α^s the same type of prior with $c_\alpha = 0.6$ and $\sigma_\alpha = 0.6$. Furthermore, in accordance with any plausible theoretical model, Baumeister and Hamilton truncate the priors for these elasticities to have respectively only negative and positive support.

[3] A more efficient algorithm can be obtained if proposals are generated as $\alpha^{\star} = \alpha^{(m-1)} + \xi(\hat{P})\nu$, where $\hat{P}$ is the Cholesky factor of the Hessian matrix of the posterior distribution obtained in the homoskedastic case and evaluated at the parameter point that maximizes the density.

[4] Wage level is measured as the seasonally adjusted real compensation per hour and employment level is the seasonally adjusted number of people on nonfarm payrolls.

In our specification of the model, we set the exact same prior on the free elements of A and use an implied Minnesota type conditional prior for the structural parameters B_i: $\mu_i = A_i\Pi_0$, with Π_0 an $n \times k$ matrix of zeros[5].

In addition, we follow the original paper and introduce the non-dogmatic restriction that the long run effect of a labour demand shock on employment is zero. We do this by adding an independent normal prior for the sum of the coefficients in B_2 associated with lagged values of Δw_t[6].

The prior distributions for the variances $\lambda_{i,1}$ are inverse-gamma with degrees of freedom $d_{i,1} = 2$ and scale $\zeta_{i,1} = \left(d_{i,1}\hat{A}\Sigma^{(1)}\hat{A}'\right)_{i,i}$, where $\hat{A}$ is the matrix that maximizes Baumeister and Hamilton's posterior and Σ is the variance/covariance matrix of AR(1) residuals computed for the first regime.

Finally, the variance ratios $\omega_{i,s}$, $s = 2,\ldots,S$, have also inverse-gamma prior with $d_{i,s} = 2$ and scale $\zeta_{i,s} = d_{i,s}\frac{\left(\hat{A}\Sigma^{(s)}\hat{A}'\right)_{i,i}}{\left(\hat{A}\Sigma^{(1)}\hat{A}'\right)_{i,i}}$.

For this application, we use the same data analyzed in the original paper, which are quarterly and cover the period 1970:Q1-2014:Q2. In the choice of regimes, we follow Brunnermeier et al. (2021) and identify December 1989 as the endpoint of a first sub-period, and we consider 2007:Q4 as the prelude of the "Great Recession" period.

The contribution provided by heteroskedasticity to the identification of the structural parameters $\beta^d$and α^s is clear from Figure 1. The histograms show the marginal posterior distributions of the demand elasticity (left panel) and the supply elasticity (right panel) obtained from the homoskedastic SVAR (BH) and from the SVAR-H model. The continuous red lines depict the prior distributions of the two parameters.

The figure shows a considerable difference between the results of the two models, SVAR-H portrays a significantly steeper aggregate labour demand and a supply that is almost flat. Despite their conclusions differ substantially, it is worth emphasizing again that the two models are identical in all features but the time variation in the variances of shocks.

The divergent implications of the two specifications are also noticeable from the cumulated IRFs depicted in Figure 2. Although the effects associated with a unit supply shock are qualitatively similar, the responses of w_t and n_t to demand shocks are dramatically revised with the introduction of heteroskedasticity. Both the wage level and employment are less responsive to demand shocks when compared to the case with constant variances.

[5]The variance V_i is a diagonal $k \times k$ matrix with the top left element equal to δ_0, and j-th main diagonal element equal to $\frac{\delta_1 \delta_2\left(1_{j\neq i}\right)}{l^{\delta_3}}\frac{\sigma_i}{\sigma_j}$, where l is the lag of the associated regressor, σ_i and σ_j are the residual variances of an AR(1) estimated via OLS for the i-th and j-th variable of the system. In our application, δ_0=100, $\delta_1 = 0.2$, δ_2=1, $\delta_3 = 2$.

[6]This prior is centered at $-\alpha^s$and has variance $R = 0.1\sigma_i$.

Figure 1: Prior vs Posteriors.

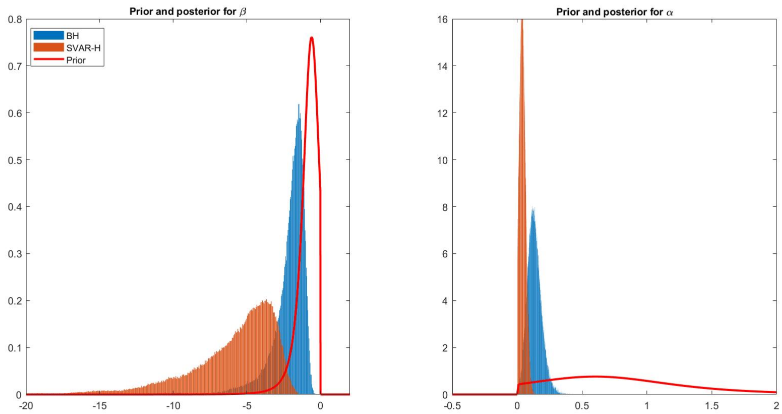

Notes: BH (blue histogram) denotes the homoskedastic SVAR with independent priors, SVAR-H (red histogram) indicates the heteroskedastic model. The red line is the prior density, which is identical in both specifications.

Figure 2: Cumulated IRFs.

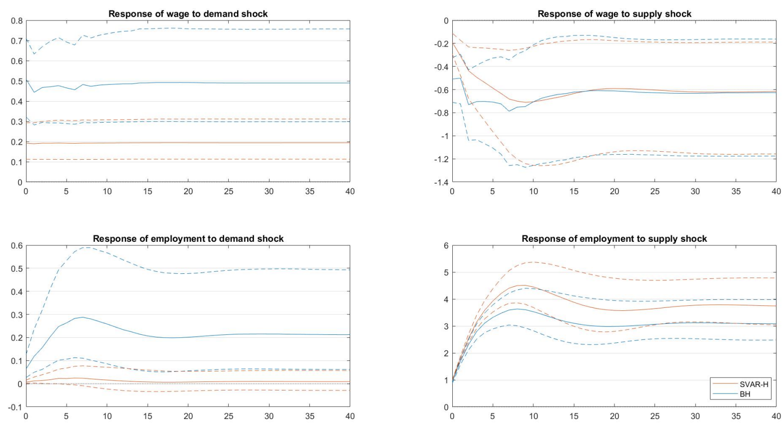

Notes: BH (blue) denotes the homoskedastic SVAR with independent priors, SVAR-H (red) indicates the heteroskedastic model. Solid lines indicate median responses while dashed lines represent 68% credible bands.

3.3 Oil Market Shocks

The second application in which we compare SVAR-H with its homoskedastic version is taken from Baumeister and Hamilton (2018) and studies the effects of different types of oil shocks on the macroeconomy.

The model is a SVAR(12) with endogenous variables $y_t = [q_t, y_t, p_t, \Delta i_t]'$, where q_t is the monthly growth rate of crude oil production, y_t is a measure of real global economic activity [7], p_t is the logarithmic difference between the acquisition cost of crude oil and the U.S. CPI, and Δi_t is the monthly change in OECD oil inventories. Baumeister and Hamilton (2018) also refine the model to take into account the fact that Δi_t contains measurement errors and represents only a share, $\chi \in [0,1]$, of the global inventories [8].

The authors argue that the SVAR with contemporaneous relations represented by:

$$\bar{A} = \begin{bmatrix} 1 & 0 & -\alpha_{qp} & 0 \\ 0 & 1 & -\alpha_{yp} & 0 \\ 1 & -\beta_{qy} & -\beta_{qp} & -\chi^{-1} \\ -\Gamma_1 & 0 & -\Gamma_3 & 1 \end{bmatrix} \tag{12}$$

appropriately describes the interactions between observable variables. However, because of the presence of measurement errors, such model is driven by innovations that are contemporaneously correlated and needs to be premultiplied by the matrix:

$$\Gamma = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & \rho & 1 \end{bmatrix} \tag{13}$$

in order to represent a system of simultaneous equations with orthogonal shocks.

In this application, the contemporaneous matrix in a SVAR representation of the type (2) is therefore $A = \Gamma\bar{A}$, and its prior distribution is the product of the priors for the individual free elements in $\bar{A}$ and the prior for ρ.

On the elasticity of oil supply, α_{qp}, the authors place a stundent-t prior with 3 degrees of freedom, location parameter $c^{\alpha}_{qp} = 0.1$ and scale $\sigma^{\alpha}_{qp} = 0.2$, truncated to be positive. On the price elasticity and the income elasticity of oil demand, β_{qp} and β_{qy}, on the other hand, they consider a stundent-t prior with the same degrees of freedom and scale, but locations at $c^{\beta}_{qp} = -0.1$ and $c^{\beta}_{qp} = 0.7$, and truncated to be respectively negative and positive. For α_{yp}, that measures the sensitiveness of global economic activity to the price of oil, they set a truncated stundent-t prior with three degrees of freedom, scale $\sigma^{\alpha}_{yp} = 0.1$ and mode $c^{\alpha}_{qp} = -0.05$.

The parameter χ represents the weight of OECD countries in the total world inventories,

[7] The measure was proposed by Kilian and Murphy (2012) and represents the cost of international shipping deflated by U.S. CPI.

[8] See section III and Appendix F of Baumeister and Hamilton (2018) for details.

and it has a Beta prior distribution with parameters $\alpha_\chi = 15$ and $\beta_\chi = 10$, to reflect the fact that the OECD group contributes for about 60% of the global oil consumption.

The priors for the remaining free elements of $\bar{A}$, Γ_1 and Γ_2, are also t-distributed with three degrees of freedom, scale equal to 0.5 and location at 0.

The parameter ρ determines the importance of measurement errors in Δi_t, and it is distributed a priori as a Beta stochastic variable, with parameters $\alpha_\rho = 3$ and $\beta_\rho = 9$, multiplied by χ.

In addition, Baumeister and Hamilton consider an asymmetric t-distribution as prior for the determinant of the matrix $\bar{A}$ and a symmetric t-distribution as prior for $h_2 = \frac{\det(\bar{A})+\alpha_{yp}\beta_{qy}}{\det(\bar{A})}$[9] to reflect the belief that the effect of a positive global activity shock should be beneficial for the world economy.

The prior for the lagged coefficients, B, is normal with Minnesota-type variances[10] and zero means, with the exception of the coefficients associated with the first lag of prices in the supply and demand equations, which have means 0.1 and -0.1 respectively.

Finally, we set inverse gamma priors for the variances $\lambda_{i,1}$ and the ratios $\omega_{i,s}$, with two degrees of freedom and scale parameters computed as described in section 3.2.

In our analysis, we use the extended version of Kilian's (2009) dataset compiled by Baumeister and Hamilton (2018). Observations are monthly and run from 1958:M1 to 2016:M12, where the pre-sample 1958:M1-1975:M1 is downweighted with a weight $\varpi = 0.5$.

We identify the first change in variances with the end of the pre-sample, and then follow Brunnermeier et al. (2021) for the definition of the remaining regimes [11].

As in the previous application, heteroskedasticity introduces in the model important identifying information. Figure 3 shows how the posteriors of the structural parameters in A differ when the variances of shocks are allowed to change across regimes.

In contrast with the homoskedastic model, the SVAR-H results support Kilian and Murphy's (2012) assumption that the price elasticity of supply is smaller than 0.0258, which has been harshly criticized by Baumeister and Hamilton. The world demand for oil is, however, much steeper in the heteroskedastic specification.

Other important differences apparent from Figure 3 are the contemporaneous coefficients in the equation describing the dynamics of inventories. According to SVAR-H, inventories have an insignificant sensitivity to the quantity of crude oil produced, while in the standard SVAR this sensitivity is negative and significant.

---

[9] The former distributions has three degrees of freedom, location $c_1 = 0.6$, scale $\sigma_1 = 1.6$ and skewness parameters $\lambda_1 = 2$. The latter distribution has three degrees of freedom, location $c_2 = 0.8$ and scale $\sigma_2 = 0.2$.

[10] Variances are $\frac{\delta_1 \delta_2 (1_{j \neq i})}{l^{\delta_3}} \frac{\sigma_i}{\sigma_j}$, with δ_0=100, $\delta_1 = 0.5$, δ_2=1, $\delta_3 = 1$.

[11] Regime 2 goes from 1975:M2 to 1979:M9 and it is characterized by the oil crisis and stagflation; regime 3 goes from 1979:M10 to 1982:M12 and defines Volcker's disinflation; regime 4 goes from 1983:M1 to 1989:M12 and includes the Saving and Loans crisis; regime 5 goes from 1990:M1 to 2007:M12 is defined as the "Great Moderation"; regime 6 goes from 2008:M1 to 2010:M12 represents the "Great Recession"; regime 7 goes from 2011:M1 to 2016:M12 and covers the recovery from the Great Recession.

Figure 3: Prior vs Posteriors.

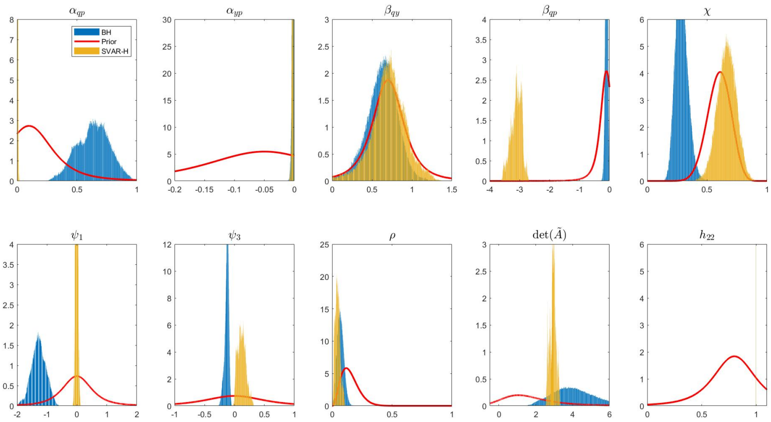

Notes: BH (blue histogram) denotes the homoskedastic SVAR with independent priors, SVAR-H (red histogram) indicates the heteroskedastic model. The red line is the prior density, which is identical in both specifications.

Finally, the two models infer different signs for the price elasticity of inventories, which is positive under SVAR-H and negative under the alternative specification.

The additional information brought in by heteroskedasticity also translates in significantly different IRFs. Figure 4 plots the responses of the four endogenous variables to oil supply shocks, global economic activity shocks, oil specific demand shock and inventories demand shocks, interpreted by Baumeister and Hamilton as speculative surprises.

One important qualitative difference between the two models results is the response of world economic activity to oil supply shocks: one unit shock that decreases oil production generates deeper and faster negative effects on global industrial production in the heteroskedastic model.

Moreover, oil specific demand shocks have less significant effects on oil production, economic activity and inventories than in the model with time invariant variances.

It is also remarkable the discrepancy in the effects of unit speculative shocks on oil production, which are positive in the homoskedastic case, but negative in SVAR-H.

Finally, under the heteroskedastic specification, inventories change less significantly in response to shocks of all types with the exception of speculative ones.

In a related paper, Braun (2021) performs a similar empirical application expoiting non-Gaussianity for identification. It is interesting to note that his results poit towars the same direction as ours do, although they are based on fairly different modelling assumptions. We interpret this as further evidence that data contains indeed precious information about structural parameters that is overlooked by homoskedastic Gaussian specifications.

Figure 4: IRFs.

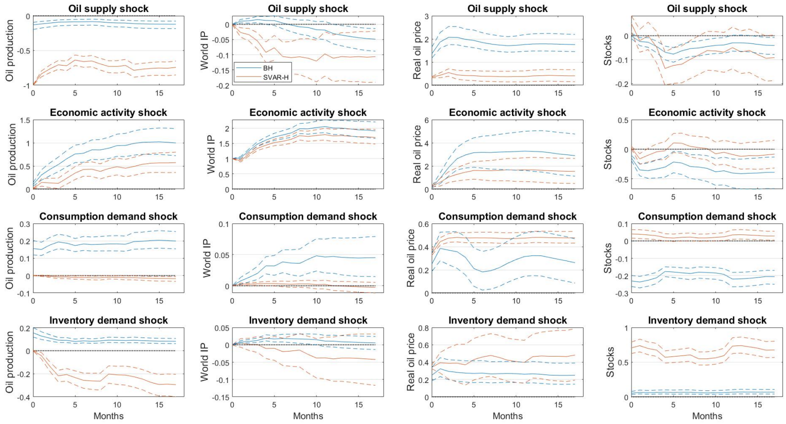

Notes: BH (blue) denotes the homoskedastic SVAR with independent priors, SVAR-H (red) indicates the heteroskedastic model. Solid lines indicate median responses while dashed lines represent 68% credible bands. Given the refinement applied to account for the presence of measurement errors in inventories, the response at horizon h of y_t to the structural shocks ε_t are computed as: $IRF_h = (\Gamma^h)_{1:n,1:n}\bar{A}^{-1}\Xi$, where Γ is the companion form of the reduced form lagged coefficients and $\Xi = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & -\chi^{-1} \\ 0 & 0 & 0 & \chi & 1 \end{bmatrix}$. See section III in Baumeister and Hamilton (2018) for details.

3.4 Monetary Policy Shocks

The last example on which we focus is the quarterly macroeconomic model of Baumeister and Hamilton (2019).

The baseline model is a SVAR(4) with a 3×1 vector of endogenous variables $y_t = [y_t, \pi_t, r_t]'$, where y_t is the output gap [12], π_t is inflation [13] and r_t is the fed fund rate observed on average during the quarter. The system dynamics is assumed to be driven by supply, demand and monetary policy shocks.

The three equations of the model are interpreted as a Phillips Curve, an aggregate demand function and a Taylor Rule, in which the contemporaneous relations are described by the matrix:

$$A = \begin{bmatrix} 1 & -\alpha^s & 0 \\ 1 & -\beta^d & -\gamma^d \\ -(1-\rho)\Gamma^y & -(1-\rho)\Gamma^\pi & 1 \end{bmatrix} \tag{14}$$

In our analysis, we adopt for the free elements in A the same prior distributions assumed in the original paper. That is, we use for Γ^y and Γ^π student-t priors with three degrees of freedom, scale parameters $\sigma_\Gamma^y = \sigma_\Gamma^\pi = 0.4$, and locations at $c_\Gamma^y = 0.5$ and $c_\Gamma^\pi = 1.5$, both truncated to be positive. The parameter ρ measures to what extent the central bank prefers smooth changes in the policy rate and it has a Beta prior distribution with mean 0.5 and standard deviation 0.2.

The priors for β^d and γ^d are both student-t distributions with three degrees of freedom,

[12]It is computed as $100\log\left(GDP_t^\star/GDP_t\right)$, where $GDP_t^\star$and GDP_t are respectively potential and actual GDP measured the Congressional Budget Office.

[13]It is computed as $100\log\left(PCE_t/PCE_{t-4}\right)$, where PCE is the PCE deflator.

scale parameter 0.4, and modes at $c_\beta = 0.75$ and $c_\gamma = -1$, but only the latter is truncated to be negative.

Finally, α^s is also t-distributed a priori with three degrees of freedom, scale $\sigma_\alpha = 0.4$ and location parameter $c_\alpha = 2$.

Additionally, we follow Baumeister and Hamilton and set asymmetric student-t priors for the combinations of parameter $h_1 = \beta^d + \gamma^d(1-\rho)\Gamma^\pi$ and $h_2 = \frac{\alpha^s\gamma^d}{\alpha^s-\beta^d}$ [14]. The role of these asymmetric priors is to put more mass on parameters values that imply negative impacts of supply shock on the inflation rate, and of monetary policy shock on economic activity.

On the lagged coeffiecients B, we set a normal distribution that implies a stationary version of the Minnesota prior for the reduced form parameters Π. In particular, we set the prior mean of B_i, conditional on A_i, to be $\mu_i' = A_i\Pi_0$, where Π_0 is 0.75 times an identity matrix, while the variances are computed as in footnote 9.

Finally, the priors for $\lambda_{i,1}$ and $\omega_{i,s}$ are inverse gammas with two degrees of freedom and scale parameters calculated as in the previous applications.

To update these priors, we use the same dataset analyzed by Baumeister and Hamilton (2019) and extend it to cover the period 1971:Q1-2019:Q4. As in the previous sections, we follow Brunnermeier et al. (2021) to determine the endpoints of variance regimes, and identify regimes changes in 1979:Q3, 1989:Q4, 2007:Q4 and 2019:Q4.

Figure 5 compares the posteriors obtained for the structural parameters in A under the SVAR-H specification and the homoskedastic alternative.

Most of the distributions are significantly affected by the introduction of heteroskedasticity. In most cases the posteriors are moved further away from the prior, but in the cases of β^d and Γ^y they are actually closer to the prior distribution. It is worth noting that this does not mean that heteroskedasticity adds little information, but, more generally, it means that the information brought in by the changes in variances is consistent with the prior specified.

The distribution that is revised more heavily is the one for the supply elasticity, which reflects a much steeper aggregate supply under SVAR-H.

The sign of the Phillips Curve coefficient β^d is reversed by both specifications with respect to the prior, but the magnitude is smaller in the heteroskedastic case. We also notice that the uncertainty about the Phillips Curve parameters is made lower by the introduction of heteroskedasticity.

In a more realistic version of the baseline model, Baumaister and Hamilton augment the three-variate system with the spread between the yield on Baa corporate bonds and the 10-year Treasury yield, the growth rate in commodity prices, and hourly wages[15].

The resulting additional parameters in the augmented contemporaneous matrix, A^*, are given student-t priors with three degrees of freedom, scale 1 and location at 0. It is also added

[14] These distributions have three degrees of freedom, scales $\sigma_{h_1} = 1$ and $\sigma_{h_2} = 0.5$, locations at $c_{h_1} = -0.1$ and $c_{h_2} = -0.3$, and skewness parameters $\lambda_{h_1} = -4$ and $\lambda_{h_1} = -2$.

[15] The growth rate in commodity prices is computed as 100 times the yearly log-difference in the CRB commodity price index. Hourly wages are the average hourly earnings of production and nonsupervisory employees.

Figure 5: Prior vs Posteriors.

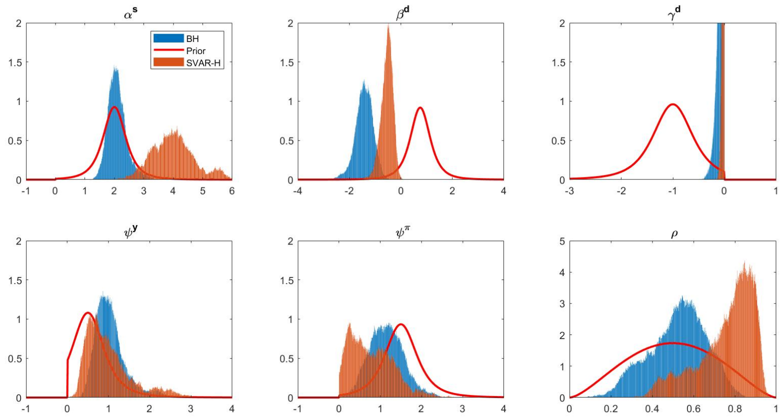

Notes: BH (blue histogram) denotes the homoskedastic SVAR with independent priors, SVAR-H (red histogram) indicates the heteroskedastic model. The red line is the prior density, which is identical in both specifications.

an independent asymmetric t prior on each of the elements of $(A^*)^{-1}$ that determine the impact response of the six variables to a monetary policy shock. These priors are such that more probability mass is assigned to combinations of structural parameters that imply an increase of the fed fund rate and the spread in response to a monetary policy shock, and a decrease in all other variables.

Figure 6 shows the responses of y_t, π_t, and r_t to supply, demand and monetary policy shocks implied by the enlarged SVAR.

The responses of the three variables to supply and demand shocks differ significantly between the homoskedastic and the heteroskedastic specifications. A unit supply shocks that decreases output produces more prolonged disinflationary effects in the SVAR-H, while the response of interest rates are not significant in the short run and becomes negative in the longer period.

Demand shocks, on the other hand, have persistent positive effects on the fed funds rate and inflation in the heteroskedastic model, while they have short-lived impacts on the fed fund rate and negative effects on inflation under the alternative specification.

Figure 7 focuses on the response of all the six variables to a unit monetary policy shock.

A monetary contraction has more recessive and persistent consequences according to SVAR-H, while the effects on the yields spread are qualitatively similar for the two specifications.

A price puzzle is present in the response of the broad measure of prices growth, but commodity inflation drops on impact for both models.

An important difference is finally indicated by the bottom-right panel of the figure. While the homoskedastic model does not detect any relevant response of wages to monetary contractions, the SVAR-H reveals a significant drop of hourly earnings after two years from the shock.

Figure 6: IRFs.

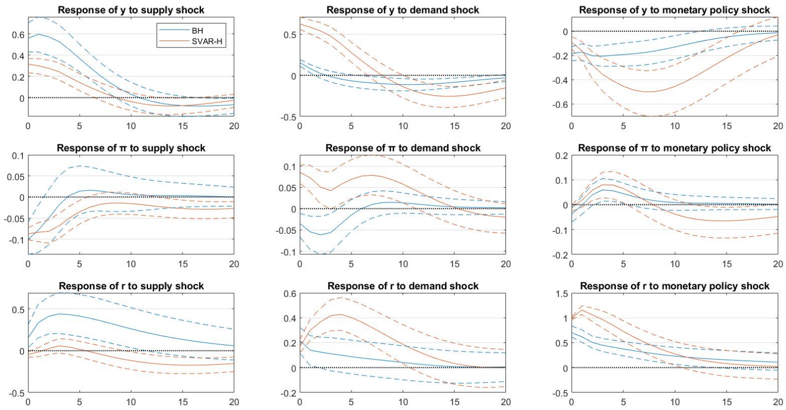

Notes: BH (blue) denotes the homoskedastic SVAR with independent priors, SVAR-H (red) indicates the heteroskedastic model. Solid lines indicate median responses while dashed lines represent 68% credible bands.

Figure 7: IRFs to monetary policy shocks.

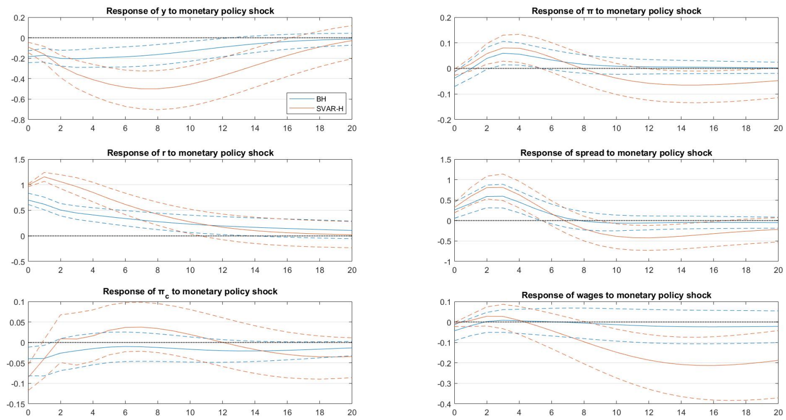

Notes: BH (blue) denotes the homoskedastic SVAR with independent priors, SVAR-H (red) indicates the heteroskedastic model. Solid lines indicate median responses while dashed lines represent 68% credible bands.

4 Proxy-SVAR and Heteroskedasticity

A Proxy-SVAR is a Structural VAR identified by means of external instruments that are correlated with the shocks of interest but uncorrelated with all other shocks. That is, given the model in equation (2), a $r \times 1$ vector of variables, $\mathbf{z}_t$, is available, and evolves according to the following process:

$$\mathbf{z}_t = \Gamma_c + \Gamma_x x_{t-1} + \Gamma_{z,1}\mathbf{z}_{t-1} + \ldots + \Gamma_{z,p}\mathbf{z}_{t-p} + \Phi\varepsilon_t + e_t \tag{15}$$

where Γ_c is a $r \times 1$ vector of intercepts, Γ_x has dimension $r \times k$, and $\Gamma_{z,l}$, $l = 1, ..., p$, are $r \times r$ matrices of lagged coefficients. The matrix Φ links the external instruments to the identified structural shocks of the SVAR and can be partitioned as $\Phi = \begin{bmatrix} \Phi_r & \mathbf{0}_{n-r} \end{bmatrix}$, in which Φ_r is a full rank $r \times r$ matrix and $\mathbf{0}_{n-r}$ is a $r \times (n-r)$ matrix of zeros. This partition makes clear why the external instruments are able to distinguish the first r shocks from all the remaining innovations. Finally, e_t is a $r \times 1$ i.i.d. vector of measurement errors that are normally distributed with mean zero and variance Σ_e, and independent from the structural shocks ε_t.

In most practical cases, and in our application in section 4.2, instruments are such that Γ_c, Γ_x and $\Gamma_{z,l}$ are all zeros, therefore equation (15) simplifies to: $\mathbf{z}_t = \Phi\varepsilon_t + e_t$.

The Proxy-SVAR can then be written compactly in reduced form as:

$$\underbrace{\begin{bmatrix} \mathbf{z}_t \\ y_t \end{bmatrix}}_{\widetilde{y}_t} = \underbrace{\begin{bmatrix} 0 \\ \Pi \end{bmatrix}}_{\widetilde{\Pi}} \underbrace{x_{t-1}}_{x_{t-1}} + \underbrace{\begin{bmatrix} \Omega_e & \Phi \\ 0 & A^{-1} \end{bmatrix}}_{\widetilde{\mathbf{A}}^{-1}} \underbrace{\begin{bmatrix} e_t \\ \varepsilon_t \end{bmatrix}}_{\widetilde{\varepsilon}_t} \tag{16}$$

where $\widetilde{\varepsilon}_t$ is normally distributed with mean zero and diagonal variance $\widetilde{\Lambda} = \begin{bmatrix} \Lambda_e & 0 \\ 0 & \Lambda \end{bmatrix}$, and Ω_e is a potentially full matrix with ones on the main diagonal, such that $\Sigma_e = \Omega_e \Lambda_e \Omega_e'$.

When r is greater than one, the Proxy-SVAR is only set identified, unless more constraints are introduced either in Φ or A^{-1}, or in both. Such constraints come often in the form of zero restrictions or sign restrictions. In the latter case, however, the identified set is only made smaller, but it is usually not reduced to a single parameter point.

In the next subsections we provide a method to point identify a Proxy-SVAR with $r > 1$ instruments and r shocks of interest, without imposing dubious restrictions on Φ or A^{-1}. Following our discussion of heteroskedasticity as a complementary identification strategy, we propose to reduce the identified set in Proxy-SVARs taking advantage of the time variation in variances of structural shocks.

4.1 Proxy-SVAR-H

Extending the reduced form Proxy-SVAR in equation (16) to account for heteroskedasticity is straightforward, it suffices to consider ε_t as generated by a normal distribution with mean zero and variance Λ_s, for $s = 1,...,S$, where S continues to indicate the number of regimes. Without loss of generality, we also set the diagonal elements of A^{-1} to 1 to fix the scale.

The likelihood of the resulting Proxy-SVAR-H is therefore:

$$p\left(y_{1:T}|\widetilde{\mathbf{A}}^{-1},\widetilde{\Pi},\widetilde{\Lambda}_{1:S}\right) = (2\pi)^{-Tn/2}\left|\det\left(\widetilde{\mathbf{A}}^{-1}\right)\right|^{-T}\left(\prod_{j=1}^{r}\lambda_{e,j}^{-T/2}\right)\left(\prod_{i=1}^{n}\lambda_{i,1}^{-T/2}\right)\left(\prod_{i=r+1}^{r+n}\prod_{s=1}^{S}\widetilde{\omega}_{i,s}^{-T_s/2}\right)\times$$

$$\exp\left\{-\frac{1}{2}\left[\sum_{s=1}^{S}\left(vec\left(\widetilde{y}_t^{(s)'}\right)-\left(I_n\otimes x^{(s)'}\right)\widetilde{\pi}\right)'\left(\left(\widetilde{\mathbf{A}}^{-1}\widetilde{\Lambda}_s\left(\widetilde{\mathbf{A}}^{-1}\right)'\right)\otimes I_{T_s}\right)^{-1}\left(vec\left(\widetilde{y}_t^{(s)'}\right)-\left(I_n\otimes x^{(s)'}\right)\widetilde{\pi}\right)\right]\right\} \tag{17}$$

Where $\widetilde{\pi} = vec\left(\widetilde{\Pi}\right)$.

The parameters in $\widetilde{\pi}$ are reduced form VAR coefficients. Therefore, given an independent Gaussian prior, $N(\widetilde{\pi}_0, V)$, conditional on $\widetilde{\mathbf{A}}^{-1}$ and $\widetilde{\Lambda}_{1:S}$, the posterior is the familiar normal distribution with variance and mean:

$$\bar{V} = \left[V^{-1}+\sum_{s=1}^{S}\left(\widetilde{\mathbf{A}}\otimes x^{(s)'}\right)'\left(\widetilde{\Lambda}_s\otimes I_{T_s}\right)^{-1}\left(\widetilde{\mathbf{A}}\otimes x^{(s)'}\right)\right]^{-1}$$

$$\bar{\pi} = \bar{V}^{-1}\left[V^{-1}\widetilde{\pi}_0+\sum_{s=1}^{S}\left(\widetilde{\mathbf{A}}\otimes x^{(s)'}\right)'\left(\widetilde{\Lambda}_s\otimes I_{T_s}\right)^{-1}vec\left(\widetilde{y}_t^{(s)'}\widetilde{\mathbf{A}}'\right)\right] \tag{18}$$

Furthermore, conditional on $\widetilde{\mathbf{A}}$, the equations of the SVAR are independent. Accordingly, given an inverse gamma prior, $IG(d_{i,1}, \zeta_{i,1})$, $i = 1,...,r+n$, for each element of $\widetilde{\Lambda}_1$, the posteriors are conjugate with $\bar{d}_{i,1} = d_{i,1}+T$ degrees of freedom, and scale parameters:

$$\bar{\zeta}_{i,1} = \zeta_{i,1}+\sum_{s=1}^{S}\omega_{i,s}^{-1}\left[\widetilde{\mathrm{A}}_i\left(y^{(s)}-\widetilde{\Pi}x^{(\mathbf{s})}\right)\left(y^{(s)}-\widetilde{\Pi}x^{(\mathbf{s})}\right)'\widetilde{\mathrm{A}}_i'\right] \tag{19}$$

Since in many applications we do not have reasons to believe that the measurement error in proxies is heteroskedastic, we focus on the case in which the first r diagonal elements of $\widetilde{\Lambda}_s$ are time invariant, that is, we set $\widetilde{\omega}_{i,s} = 1$ for $i = 1,...,r$ and $s = 2,...,S$.

As a result, we consider regime-switching variance ratios only for the structural shocks ε_t. For given priors, $\widetilde{\omega}_{i,s} \sim IG(d_{i,s}, \zeta_{i,s})$, $i = r+1,...,r+n$ and $s = 2,...,S$, the posteriors are then also inverse gammas with moments:

$$\bar{d}_{i,s} = d_{i,s}+T_s$$

$$\bar{\zeta}_{i,s} = \zeta_{i,s}+\widetilde{\lambda}_{i,1}^{-1}\left[\widetilde{\mathrm{A}}_i\left(y^{(s)}-\widetilde{\Pi}x^{(\mathbf{s})}\right)\left(y^{(s)}-\widetilde{\Pi}x^{(\mathbf{s})}\right)'\widetilde{\mathrm{A}}_i'\right] \tag{20}$$

Finally, the conditional posterior of $\widetilde{\mathbf{A}}^{-1}$ is not of a standard form because of the term

$\left|\det\left(\widetilde{\mathbf{A}}^{-1}\right)\right|^{-T}$ appearing in the likelihood. Provided a prior for the n_α free elements in $\widetilde{\mathbf{A}}^{-1}$, $p(\alpha)$, we can thus write its posterior distribution as:

$$p\left(\alpha|Y,\widetilde{\Pi},\widetilde{\Lambda}_{1:S}\right) \propto p(\alpha)\left|\det\left(\widetilde{\mathbf{A}}^{-1}\right)\right|^{-T} \times$$

$$\exp\left\{-\frac{1}{2}\left[\sum_{s=1}^{S}\left(vec\left(\widetilde{y}_t^{(s)'}\right)-\left(I_n\otimes x^{(s)'}\right)\widetilde{\pi}\right)'\left(\left(\widetilde{\mathbf{A}}^{-1}\widetilde{\Lambda}_s\left(\widetilde{\mathbf{A}}^{-1}\right)'\right)\otimes I_{T_s}\right)^{-1}\left(vec\left(\widetilde{y}_t^{(s)'}\right)-\left(I_n\otimes x^{(s)'}\right)\widetilde{\pi}\right)\right]\right\} \tag{21}$$

and use a Metropolis-Hastings step to approximate it.

The MCMC algorithm we use to draw from the posterior of the Proxy-SVAR-H parameter can then be summarized as:

1. Draw from $p\left(\widetilde{\Lambda}_s|Y,\left(\widetilde{\mathbf{A}}^{-1}\right)^{(m-1)},\widetilde{\Pi}^{(m-1)}\right)$ for $s=1,...,S$, which is accomplished by:
   (a) drawing from $p\left(\widetilde{\lambda}_{i,1}|Y,\left(\widetilde{\mathbf{A}}^{-1}\right)^{(m-1)},\widetilde{\Pi}^{(m-1)},\widetilde{\omega}_{i,2:S}^{(m-1)}\right)$ for $i=1,...,r+n$;
   (b) drawing from $p\left(\widetilde{\omega}_{i,s}|Y,\left(\widetilde{\mathbf{A}}^{-1}\right)^{(m-1)},\widetilde{\Pi}^{(m-1)},\widetilde{\lambda}_{i,1}^{(m)}\right)$ for $i=r+1,...,n$ and for $s=2,...,S$;
2. Draw from an approximation of $p\left(\alpha|Y,\widetilde{\Pi}^{(m-1)},\widetilde{\Lambda}_{1:S}^{(m)}\right)$ using a Random Walk Metropolis-Hastings algorithm, which:
   (a) draws a candidate $\alpha^\star=\alpha^{(m-1)}+\xi\nu$, where $\nu\sim t(2)$ and has dimension $n_\alpha\times 1$, while ξ is a tuning parameter that guarantees an acceptance rate between 30% and 40%;
   (b) compute $\vartheta=\min\left\{\frac{p\left(\alpha^\star|Y,\widetilde{\Pi}^{(m-1)},\widetilde{\Lambda}_{1:S}^{(m)}\right)}{p\left(\alpha^{(m-1)}|Y,\widetilde{\Pi}^{(m-1)},\widetilde{\Lambda}_{1:S}^{(m)}\right)},1\right\}$, and set $\alpha^{(m)}=\alpha^\star$ with probability ϑ and $\alpha^{(m)}=\alpha^{(m-1)}$ with probability $(1-\vartheta)$;
3. Draw from $p\left(\widetilde{\pi}^{(m-1)}|Y,\left(\widetilde{\mathbf{A}}^{-1}\right)^{(m)},\widetilde{\Lambda}_{1:S}^{(m)}\right)$.

Which is iterated for $m=1,...,M$, where M is a sufficiently large number, discarding the first M_0 draws[16].

Notice that, since the coefficients in the first r rows of $\widetilde{\Pi}$ are all zeros and the last n reduced form errors are independent from the measurement errors, the last step of the sampler can be performed by focusing only on the last n equations of the augmented VAR.

[16] In our empirical application $M=2,000,000$ and $M_0=1,000,000$.

4.2 Mertens and Ravn (2013)

One of the most important studies exploiting the availability of external instruments in a SVAR framework is Mertens and Ravn (2013). The authors construct two narrative measures correlated with personal income tax (APITR) and corporate income tax (ACITR) shocks for the U.S.. They then acknowledge the fact that these proxies contain measurement errors and therefore rely on a Proxy-SVAR analysis.

This is a typical example of the situation in which $r > 1$ ($r = 2$ in this case) instruments are available to identify $r > 1$ shocks. As a result, the structural model is only set-identified.

Mertens and Ravn (2013) obtain point identification by imposing that the matrix Φ_r is lower triangular, ordering either the APITR or the ACITR instrument first. In other words, the authors assume that the instrument ordered first only contains information about one shock, while the other is correlated with both fiscal shocks.

As argued by Giacomini et al. (2020), this additional restriction may be undesirable because it introduces a significant degree of arbitrariness in the identification scheme. To avoid this problem, Giacomini et al. (2020) propose to relax the zero restriction and to base inference on the whole identified set.

Our strategy is instead to substitute the zero restriction with heteroskedasticity, basing inference on a point identified model.

In this section we revisit Mertens and Ravn's (2013) results through the lenses of a Proxy-SVAR-H and show that the information brought in by heteroskedasticity can be substantial.

The SVAR in the original paper has 4 lags and a $n \times 1$ vector of endogenous variables, $y_t = [tp_t, tc_t, bp_t, bc_t, g_t, y_t, d_t]'$, where tp_t and tc_t are respectively the real personal and corporate average income tax rate, bp_t and bc_t are the real personal and corporate income tax base per capita, g_t measures government purchases of final goods per capita, y_t is the real GDP per capita, and d_t is the real Federal Government debt per capita.

Data are quarterly and cover the period 1951:Q1-2006:Q4. We define a first variance regime that goes from 1951:Q1 to 1971:Q1 and can be called the "Bretton Woods" regime, and select 1996:Q4 as a further break point separating a period of deficit increase from the subsequent period of deficit cuts.

The prior we set for the lagged coefficients, $\widetilde{\Pi}$, are Gaussian with an identity matrix as mean and Minnesota-type variances[17]. The inverse gamma priors for the diagonal elements of $\widetilde{\Lambda}_1$ and for the variance ratios $\omega_{i,s}$ have two degrees of freedom and scale parameters computed as in section 3.2.

Finally, all the free elements in $\widetilde{\mathbf{A}}^{-1}$ have independent student-t prior distributions with three degrees of freedom, scale 0.5 and location at 0.

Figures 8 and 9 depict the responses to one-unit tax shock for both the homoskedastic and the heteroskedastic version of the Proxy-SVAR. The IRFs in the first figure are generated by

[17] Variances are $\frac{\delta_1 \delta_2 (1_{j \neq i})}{l^{\delta_3}} \frac{\sigma_i}{\sigma_j}$, with δ_0=100, $\delta_1 = 0.5$, δ_2=1, $\delta_3 = 2$.

Figure 8: IRFs to APTIR shocks.

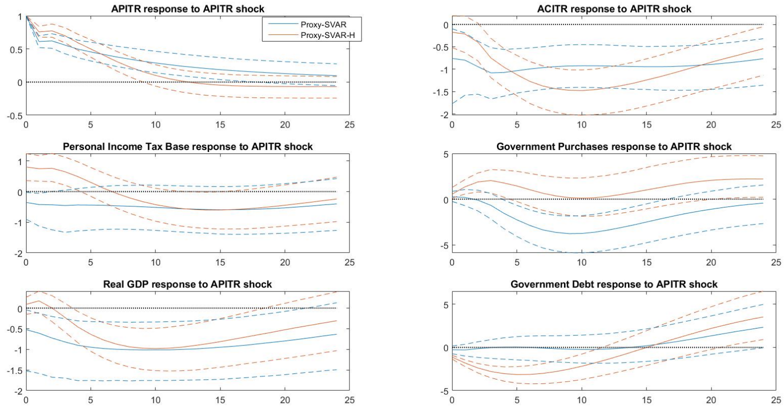

Notes: Proxy-SVAR (blue) denotes the homoskedastic SVAR with independent priors, Proxy-SVAR-H (red) indicates the heteroskedastic model. Solid lines indicate median responses while dashed lines represent 68% credible bands.

a fiscal surprise that rises APITR by 1 percentage point, while the second figure plots the responses to a 1% unexpected increase in ACITR.

It is clear that heteroskedasticity adds important information in many directions. The Proxy-SVAR-H model makes evident that both types of fiscal contractions have deep negative effects on output.

Another important feature highlighted by the heteroskedastic model is that, following a corporate income tax shock, corporate income tax base actually decreases. This mechanism, together with the negative effect on output, leads government debt to rise in response to the unexpected corporate tax increase.

On the other hand, a personal income tax shock raises the tax base and allows a reduction of the government debt as one would expect.

Finally. government purchases increase after both types of fiscal surprises, but this effect tends to be reabsorbed within one year from the shock.

5 Conclusions

In this paper we propose to combine heteroskedasticity with sign restrictions and external instruments in order to sharpen identification. Introducing regime shifting in the variances of structural shocks, we allow the data to be informative about which of the structural representations compatible with sign restrictions or with conditions coming from external instruments is admissible.

On the one hand, we overcome the limitations of the identification strategies based on sign restrictions or external instruments, i.e. potentially large identified sets, taking advantage of the information introduced by heteroskedasticity. On the other hand, we overcome the labeling

Figure 9: IRFs to ACTIR shocks.

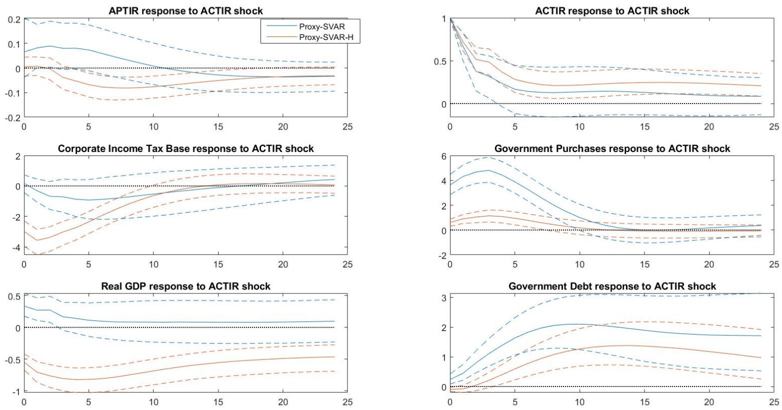

Notes: Proxy-SVAR (blue) denotes the homoskedastic SVAR with independent priors, Proxy-SVAR-H (red) indicates the heteroskedastic model. Solid lines indicate median responses while dashed lines represent 68% credible bands.

problem inherent in identification through heteroskedasticity, considering additional identifying information in the form of sign restrictions or external instruments.

We illustrate the convenience of this hybrid approach by revisiting empirical results presented by seminal papers that are based on either sign restrictions or proxy variables. In particular, we reconsider the applications in Baumeister and Hamilton (2015, 2018, 2019), which analyze the effects of labour market, oil market and monetary policy shocks respectively. Moreover, we reproduce the Proxy-SVAR model in Mertens and Ravn (2013), whose aim is to assess the effects of Average Personal Income Tax Rate and Average Corporate Income Tax Rate shocks, and we do this by relaxing the zero restrictions that allows them to achieve point identification.

In all the empirical applications, allowing for heteroskedasticity does sharpen identification, and leads to economically relevant changes in the effects of the identified structural shocks.

References

[1] Baumeister, C. and Hamilton, J.D. , 2015. “Sign restrictions, structural vector autoregressions, and useful prior information”, *Econometrica.* 83 (5), 1963–1999;

[2] Baumeister, C. and Hamilton, J.D. , 2018. “Inference in Structural Vector Autoregressions When the Identifying Assumptions are Not Fully Believed: Re-evaluating the Role of Monetary Policy in Economic Fluctuation”, *Journal of Monetary Economics*, 100, 48-65;

[3] Baumeister, C. and Hamilton, J.D. , 2019. “Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks”, *American Economic Review*, 109(5), 1873-1910;

[4] Braun, R., 2021, “The importance of supply and demand for oil prices: evidence from a SVAR identified by non-Gaussianity”, Manuscript;

[5] Braun, R. and Brüggemann, R., 2020. "Identification of SVAR Models by Combining Sign Restrictions With External Instruments," *Working Paper Series of the Department of Economics*, University of Konstanz 2020-01, Department of Economics, University of Konstanz;

[6] Brunnermeier, M., Darius, P., Sastry, K. A. and Sims, C. A., 2021, "Feedbacks: Financial Markets and Economic Activity." *American Economic Review*, 111 (6): 1845-79;

[7] Canova, F., and De Nicolo, G.. 2002. “Monetary Disturbances Matter for Business Fluctuations in the G-7.” *Journal of Monetary Economics* 49 (6): 1131–59;

[8] Faust, J. 1998. “The Robustness of Identified VAR Conclusions about Money.” *International Finance Discussion Paper* 610;

[9] Giacomini, R., Kitagawa, T. and Read, M., 2020. “Robust Bayesian Inference in Proxy Svars”. *CEPR Discussion Paper* No. DP14626, Available at SSRN: https://ssrn.com/abstract=3594257;

[10] Kilian, L., 2009. “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market.” *American Economic Review* 99 (3): 1053–69;

[11] Kilian, L. , Lütkepohl, H. , 2017. “Structural Vector Autoregressive Analysis”. Cambridge University Press, Cambridge;

[12] Kilian, L., and Murphy, D. P., 2012. “Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models.” *Journal of the European Economic Association* 10 (5): 1166–88.

[13] Lanne, M., H. Lütkepohl, and K. Maciejowska, 2010, “Structural Vector Autoregressions with Markov Switching,” *Journal of Economic Dynamics and Control*, 34: 121-31;

[14] Lewis, D. J., 2018a. "Robust inference in models identified via heteroskedasticity," *Staff Reports 876*, Federal Reserve Bank of New York;

[15] Daniel J. Lewis, 2018b. "Identifying shocks via time-varying volatility," *Staff Reports 871*, Federal Reserve Bank of New York;

[16] Lakdawala, A., 2019. “Decomposing the effects of monetary policy using an external Instruments SVAR.” *Journal of Applied Econometrics*, 34 (6), 934–950;

[17] Lütkepohl H., Woźniak T., 2020, “Bayesian inference for structural vector autoregressions identified by Markov-switching heteroskedasticity,” *Journal of Economic Dynamics & Control* 113, 1–21;

[18] Mertens, K. and Ravn, M.O., 2013. “The Dynamic Eects of Personal and Corporate Income Tax Changes in the United States”. *American Economic Review* 103 (4), pp. 1212-47;

[19] Piffer, M. and Podstawski, M., (2018). “Identifying Uncertainty Shocks Using the Price

of Gold". *The Economic Journal* 128 (616), pp. 3266-3284;

[20] Rigobon, R., 2003, "Identication through Heteroskedasticity", *Review of Economics and Statistics*, 85, 777-792;

[21] Rigobon, R. and Sack, B., 2004. "The Impact of Monetary Policy on Asset Prices," *Journal of Monetary Economics*, 51, 1553–1575;

[22] Uhlig, H., 2005. "What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure." *Journal of Monetary Economics* 52 (2): 381–419;

[23] Sentana, E., and Fiorentini, G.. 2001, "Identication, Estimation and Testing of Conditionally Heteroskedastic Factor Models", *Journal of Econometrics* 102, 143-164;

[24] Sims, C., 1980. "Macroeconomics and Reality." *Econometrica*, 48(1), 1-48. doi:10.2307/1912017;

[25] Sims, C. A. 2020. "SVAR Identification through Heteroskedasticity with Misspecified Regimes." Princeton University;

[26] Sims C. A. and T. Zha, 1998, "Bayesian Methods for Dynamics Multivariate Models", *International Economic Review*, 39, 949-968;

[27] Stock, J.H., 2008. "What's New in Econometrics-Time Series, Lecture 7: Structural VARs. Cambridge", MA.: National Institute for Economic Research;

[28] Stock, J.H. and Watson, M.W., 2012. "Disentangling the Channels of the 2007-09 Recession". *Brookings Papers on Economic Activity, Spring*, pp. 81-156.