Journal of Business & Economic Statistics

ISSN: 0735-0015 (Print) 1537-2707 (Online) Journal homepage: www.tandfonline.com/journals/ubes20

Avoiding Unintentionally Correlated Shocks in Proxy Vector Autoregressive Analysis

Martin Bruns, Helmut Lütkepohl & James McNeil

To cite this article: Martin Bruns, Helmut Lütkepohl & James McNeil (2025) Avoiding Unintentionally Correlated Shocks in Proxy Vector Autoregressive Analysis, Journal of Business & Economic Statistics, 43:4, 1119-1131, DOI: 10.1080/07350015.2025.2476699

To link to this article: https://doi.org/10.1080/07350015.2025.2476699

JOURNAL OF BUSINESS & ECONOMIC STATISTICS
2025, VOL. 43, NO. 4, 1119–1131
https://doi.org/10.1080/07350015.2025.2476699

Avoiding Unintentionally Correlated Shocks in Proxy Vector Autoregressive Analysis

Martin Bruns[a], Helmut Lütkepohl[b], and James McNeil[c]

[a]School of Economics, University of East Anglia, Norwich, UK; [b]DIW Berlin and Freie Universität Berlin, Berlin, Germany; [c]Department of Economics, Dalhousie University, Halifax, Nova Scotia, Canada

ABSTRACT
Noting that the shocks in vector autoregressive models can be correlated if a number of shocks is identified individually by multiple proxy variables, we propose a Generalized Method of Moments (GMM) approach for estimation that enforces uncorrelated shocks. We point out that if each proxy identifies exactly one shock and is uncorrelated with all other shocks, uncorrelatedness of the shocks provides over-identifying restrictions that can be used in our approach to improve the estimation efficiency of the structural parameters. It also opens up the possibility to use Hansen's J-test to check the model specification. Our method generalizes other GMM proposals that work under more restrictive assumptions. We illustrate its usefulness by two empirical examples from the recent literature.

ARTICLE HISTORY
Received July 2024
Accepted March 2025

KEYWORDS
Correlated shocks; External instruments; Generalized method of moments; Structural vector autoregression; Proxy VAR

1. Introduction

In recent years it has become increasingly popular to identify structural shocks in structural vector autoregressive (VAR) analysis by external instruments or proxies. In a number of studies, more than one shock of interest is identified in this way. For example, Hou (2024) analyses the response of a range of U.S. macroeconomic variables to five shocks—an oil shock, a monetary policy shock, a productivity shock, a liquidity and financial risk shock, and a fiscal policy shock—which he identifies using a set of proxies. It is well known that while multiple proxies can identify linear combinations of structural shocks, further information is required to disentangle the individual shocks of interest (see Hou (2024), and the citations therein). In some studies, it is assumed that a single proxy or a set of proxies is correlated with a single shock only and is uncorrelated with all other shocks and the shocks are identified one-by-one (e.g., Altavilla et al. 2019). However, it was pointed out by Stock and Watson (2012) and emphasized more recently by other researchers, including Gregory, McNeil, and Smith (2024), that in this case the shocks may not be instantaneously uncorrelated any more, violating a standard assumption of structural VAR analysis. This can happen even if the proxies are mutually uncorrelated and individually satisfy the standard relevance and exogeneity conditions that the proxy VAR literature typically assumes for valid proxies. Having uncorrelated structural shocks is, for example, important for the proper interpretation of impulse responses and for performing forecast error variance decompositions (FEVDs) (see also the related discussions in Ramey (2016), Stock and Watson (2018), and Gregory, McNeil, and Smith (2024)). Thus, it may be problematic that some of the shocks considered in recent structural VAR analyses are not instantaneously uncorrelated.

Gregory, McNeil, and Smith (2024) present a generalized method of moments (GMM) method that provides uncorrelated shocks if the condition of each proxy being correlated with a single shock only is satisfied. Their method requires that, in a K-dimensional VAR, at least $K-1$ proxies are available so that all K shocks are identified. Unfortunately, that condition is not satisfied in a range of proxy VAR studies (e.g., Lunsford 2015; Piffer and Podstawski 2017; Lakdawala 2019; Jarociński and Karadi 2020; Känzig 2021; Fanelli and Marsi 2022). In these papers typically just-identifying or set-identifying restrictions are used for the structural parameters. Such restrictions cannot be checked by statistical tests and, thus, have to be backed by convincing subject matter arguments.

In this article, we propose a simple GMM method that works more generally even when only a subset of the shocks are identified and it can also take advantage of over-identifying restrictions. The method focuses on the structural parameters of interest in the GMM objective function and, hence, will typically result in a computationally simple optimization problem. If each proxy is only correlated with a single shock, it is shown that uncorrelatedness of the shocks provides over-identified structural parameters. A device proposed by Crepon, Kramarz, and Trognon (1997) is used to ensure that in that case our method also provides asymptotically efficient estimators under standard assumptions. Using this device requires the optimal weighting matrix in the GMM objective function for which we derive a closed-form expression. It also ensures a valid J-test for model misspecification if the moment conditions over-identify

CONTACT Martin Bruns martin.bruns@uea.ac.uk School of Economics, University of East Anglia, Norwich Research Park, NR4 7TJ, Norwich, UK.
Supplementary materials for this article are available online. Please go to www.tandfonline.com/UBES.

the structural parameters and opens up the use of this tool in a proxy VAR context for assessing the identifying assumptions. In our proposed setup, the J-test is focussed on the structural parameters of interest instead of applying it to the whole set of moment conditions related to all the parameters. Thereby the test has good power even in small samples against alternatives relevant in a proxy VAR analysis.

We will present examples of empirical studies, where using the proxies one-by-one leads to correlated shocks while our GMM approach fixes the problem and makes a difference for the implied impulse responses and FEVDs. Thereby we show that avoiding correlated shocks is not only a theoretical problem but is relevant for applied work. In these examples we also show that our J-test is a useful tool for model diagnostics in proxy VAR studies.

Our approach is related to Hou (2024), who proposes a Bayesian method that can be used when only a subset of shocks are identified and ensures the identified shocks are uncorrelated. Over-identification conditions can then be tested by comparing the marginal likelihoods of the proxy-SVAR with and without the over-identifying restrictions imposed. Our GMM approach complements this existing work by introducing a method which is applicable in similar contexts but lies instead within the frequentist paradigm.

The remainder of the article is organized as follows. In the next section we present the model setup formally and discuss the GMM procedure that can be used for solving the problem of getting correlated shocks. In Section 3, we present simulation results showing that our proposed methods work in small samples and in Section 4 we consider empirical examples. Section 5 concludes. Some mathematical derivations as well as additional results related to the simulations and empirical examples are provided in an online supplement.

The following general notation will be used throughout. The operator $\text{vec}(\cdot)$ is the usual column vectorization operator for a matrix. $\text{vech}(\cdot)$ is the corresponding operator vectorizing a square matrix from the main diagonal downwards, and $\text{vh}(\cdot)$ is the vectorization operator that collects only the elements below the main diagonal of a matrix in a vector. Moreover, we use the $(\frac{1}{2}m(m-1)\times m^2)$ selection matrix $\mathbf{S}_m$ that selects the elements below the main diagonal of an $(m\times m)$ matrix M from $\text{vec}(M)$, that is, $\text{vh}(M)=\mathbf{S}_m\text{vec}(M)$.

2. Model Setup, Estimation, and Inference

2.1. The Model

Our basic model is a K-dimensional reduced-form VAR process,

$$\begin{aligned} y_t &= \nu + A_1 y_{t-1} + \cdots + A_p y_{t-p} + u_t \\ &= (\nu, A_1, \ldots, A_p) Y_{t-1} + u_t, \end{aligned} \tag{1}$$

where u_t is a zero mean white noise process with nonsingular covariance matrix Σ_u, that is, $u_t \sim (0,\Sigma_u)$ and $Y_{t-1} = (1, y'_{t-1}, \ldots, y'_{t-p})'$ is a $(Kp+1)$-dimensional column vector. This model is assumed to be the data generating process (DGP).

The vector of structural shocks is denoted by $\mathbf{w}_t = (w_{1t},\ldots,w_{Kt})'$. It is obtained from the reduced-form errors, u_t, by a linear transformation, $\mathbf{w}_t = B^{-1}u_t$. The $(K\times K)$ matrix $B=[b_{ij}]$ contains the impact effects of the structural shocks and $B\Sigma_{\mathbf{w}}B' = \Sigma_u$, where $\Sigma_{\mathbf{w}}$ is the covariance matrix of $\mathbf{w}_t$. As in much of the structural VAR literature, the shocks are assumed to be instantaneously uncorrelated and, hence, the transformation matrix B is such that the covariance matrix $\Sigma_{\mathbf{w}}$ is diagonal and the structural shocks $\mathbf{w}_t$ are instantaneously uncorrelated by construction.

As we are also considering partially identified models, we partition $\mathbf{w}_t$ in K_1- and $(K-K_1)$-dimensional subvectors $\mathbf{w}_{1t} = (w_{1t},\ldots,w_{K_1t})'$ and $\mathbf{w}_{2t} = (w_{K_1+1,t},\ldots,w_{Kt})'$ such that $\mathbf{w}'_t = (\mathbf{w}'_{1t},\mathbf{w}'_{2t})$. The first K_1 shocks, $\mathbf{w}_{1t}$, are the structural shocks of interest. They have to be identified properly, while $\mathbf{w}_{2t}$ contains shocks which are not in the focus of the analysis and are, hence, not necessarily identified as proper economic shocks. The matrix of impact effects, B, is partitioned accordingly as $B=[B_1 : B_2]$, where B_1 is $(K\times K_1)$ and B_2 is $(K\times(K-K_1))$. In other words, B_i contains the impact effects of the shocks $\mathbf{w}_{it}$, $i=1,2$.

The impact effects, B, are the structural parameters of the model. The shocks $\mathbf{w}_{1t}$ are identified if the B_1 matrix is identified. Having the matrix B_1, we can compute the structural impulse responses to the $\mathbf{w}_{1t}$ shocks for propagation horizon h as

$$\Theta_{1h} = \Phi_h B_1, \tag{2}$$

where the Φ_h are reduced-form quantities obtained recursively from the VAR slope coefficients as $\Phi_h = \sum_{j=1}^h \Phi_{h-j}A_j$, with $\Phi_0 = I_K$, for $h=0,1,\ldots$, and $A_j = 0$ for $j>p$ (see, e.g., Lütkepohl 2005, sec. 2.1.2).

Each column of B contains the impact effects of a single shock on all the K variables. Denoting by b_k the k-th column of B, the k-th shock can be obtained from the reduced-form residuals as

$$w_{kt} = b'_k\Sigma_u^{-1}u_t / b'_k\Sigma_u^{-1}b_k \tag{3}$$

(see, e.g., Stock and Watson (2018) and Bruns and Lütkepohl (2022, Appendix A.1)).

2.2. Identification via Proxy Variables

Identification of the structural parameters and, hence, the structural shocks is assumed to be based on a set of N instrumental variables (proxies) $z_t = (z_{1t},\ldots,z_{Nt})'$ satisfying

$$\begin{aligned} &\mathbb{E}(\mathbf{w}_{1t}z'_t) = \Sigma_{\mathbf{w}_1z} \neq 0, \quad \Sigma_{\mathbf{w}_1z}\ (K_1\times N), \\ &\text{rk}(\Sigma_{\mathbf{w}_1z}) = K_1 \quad \text{(relevance)}; && (4) \\ &\mathbb{E}(\mathbf{w}_{2t}z'_t) = 0 \quad \text{(exogeneity)}. && (5) \end{aligned}$$

These conditions imply that

$$\mathbb{E}(u_t z'_t) = B\mathbb{E}(\mathbf{w}_t z'_t) = B_1\Sigma_{\mathbf{w}_1z}. \tag{6}$$

Obviously, there must be at least as many proxies as there are identified shocks such that $N\geq K_1$, to satisfy the rank condition for $\Sigma_{\mathbf{w}_1z}$ which ensures that the N proxies contain identifying information for all shocks in $\mathbf{w}_{1t}$. As we can estimate $B_1\Sigma_{\mathbf{w}_1z}$ by the usual covariance matrix estimator

$$\overline{\hat{u}z} = \frac{1}{T}\sum_{t=1}^T \hat{u}_t z'_t, \tag{7}$$

where the $\hat{u}_t$ are reduced-form least squares (LS) residuals, the proxies contain identifying information for the first K_1

structural shocks collectively but the shocks are not necessarily individually identified. The usual techniques for identifying B_1 have been used in the related literature. For example, Mertens and Ravn (2013) place just-identifying zero restrictions on the impact effects and Jarociński and Karadi (2020) use sign restrictions to disentangle the shocks.

Alternatively, restrictions can be imposed on $\Sigma_{\mathbf{w}_1 z}$. For example, the proxies may be constructed such that each proxy is correlated with just one shock and $\Sigma_{\mathbf{w}_1 z}$ is a diagonal square matrix. In that case the shocks will be identified individually because the right-hand side of (6) will consist of multiples of the impact effects of the shocks that will provide multiples of the shocks via the relation (3). In general, the proxies may be such that zero restrictions can be imposed on $\Sigma_{\mathbf{w}_1 z}$ without being a diagonal matrix. For example, Lakdawala (2019) considers a triangular $\Sigma_{\mathbf{w}_1 z}$ matrix.

If the impact effects of the kth shock are estimated as $T^{-1}\sum_{t=1}^{T}\hat{u}_t z_{kt}$, we will refer to this approach as the *one-by-one proxy VAR approach* to estimating the impact effects of the shocks of interest. Obviously, this estimator is identical to the one obtained by using the estimator in equation (7) and thereby estimating the impact effects of all K_1 shocks of interest jointly. The shocks obtained in this way may, however, be instantaneously correlated, as pointed out recently by Gregory, McNeil, and Smith (2024), because there is no mechanism that enforces uncorrelatedness.

Gregory, McNeil, and Smith (2024) propose a GMM approach that ensures uncorrelated (orthogonal) shocks. They use moment conditions obtained from the assumption that each proxy is correlated with one shock only which implies that all the other shocks are uncorrelated with the proxy, giving $K-1$ moment conditions for each proxy. In addition, they use that all the shocks are mutually uncorrelated. Thereby they obtain $\frac{1}{2}K(K-1)$ further moment conditions. Finally, they standardize the B matrix to have a unit diagonal. In other words, they assume that each shock has a unit impact effect for one of the variables. Thereby they have to estimate just $K(K-1)$ structural parameters in the B matrix. A drawback of the Gregory, McNeil, and Smith (2024) approach is that it works only if at least $K-1$ shocks are identified by proxies. Otherwise they do not have enough moment conditions to identify all the parameters. Clearly, there are many examples of proxy VAR studies where fewer than $K-1$ shocks are identified by proxies and, hence, their approach does not work any more. It is also a disadvantage of their approach that, for each shock, they have to take a stand on a specific variable having a nonzero instantaneous response. There are many studies where the response of the variables to the shocks is not known before the analysis is conducted.

In the following, we present a GMM method that works more generally also if $K_1 < K-1$ because we use a different set of moment conditions. Our moment conditions are focussed on the first K_1 shocks of interest that are identified by proxies and do not involve moment conditions related to other shocks.

From (6) we get KN moment conditions

$$\mathbb{E}(u_t z_t' - B_1\Sigma_{\mathbf{w}_1 z}) = 0. \tag{8}$$

Moreover, using $u_t = B\mathbf{w}_t$ and, hence, $\Sigma_u = B\Sigma_{\mathbf{w}}B'$, where $\Sigma_{\mathbf{w}}$ is the covariance matrix of the $\mathbf{w}_t$ shocks, we have

$$\begin{aligned}\mathbb{E}(B_1'\Sigma_u^{-1}u_t u_t'\Sigma_u^{-1}B_1) &= B_1'B'^{-1}\Sigma_{\mathbf{w}}^{-1}B^{-1}B_1\\ &= [I_{K_1}:0]\Sigma_{\mathbf{w}}^{-1}\begin{bmatrix} I_{K_1}\\ 0\end{bmatrix}.\end{aligned} \tag{9}$$

Uncorrelated shocks imply that $\Sigma_{\mathbf{w}}$ is a diagonal matrix. Hence, considering only the elements below the main diagonal of the left-hand side matrix, we get a further set of $\frac{1}{2}K_1(K_1-1)$ moment conditions

$$\mathbb{E}[\mathrm{vh}(B_1'\Sigma_u^{-1}u_t u_t'\Sigma_u^{-1}B_1)] = 0. \tag{10}$$

Note that there are $KK_1 + K_1N$ free parameters in B_1 and $\Sigma_{\mathbf{w}_1 z}$ and we have $KN + \frac{1}{2}K_1(K_1-1)$ moment conditions in (8) and (10). Thus, in general there are fewer moment conditions than parameters to estimate. Hence, we have an under-identified estimation problem. However, it may be possible to construct the proxies such that restrictions can be imposed on $\Sigma_{\mathbf{w}_1 z}$. For example, one may construct proxies which are correlated only with a subset of the shocks in $\mathbf{w}_{1t}$. Then we can impose zero restrictions on $\Sigma_{\mathbf{w}_1 z}$.

If we have enough restrictions such that the moment conditions in (8) and (10) at least just-identify B_1 and $\Sigma_{\mathbf{w}_1 z}$ we can use GMM estimation. For example, for the case of specific interest in the following, where each proxy is correlated with a single shock only such that $N = K_1$ and $\Sigma_{\mathbf{w}_1 z}$ is a diagonal matrix, the moment conditions over-identify the parameters if $K_1 > 1$.

Note, however, that the moment conditions depend on the reduced-form VAR parameters $\alpha = \mathrm{vec}(\nu, A_1, \dots, A_p)$ via $u_t(\alpha) = y_t - (Y_{t-1}' \otimes I_K)\alpha$ and $\sigma = \mathrm{vech}(\Sigma_u)$. Thus, if we focus on estimating $\beta = \mathrm{vec}(B_1)$ and the unrestricted elements of $\Sigma_{\mathbf{w}_1 z}$, which we collect in the vector δ, we still have to account for nuisance parameters $\gamma = (\alpha', \sigma')'$. The standard moment conditions for LS estimation of the reduced-form VAR parameters are

$$\mathbb{E}[(Y_{t-1}\otimes I_K)y_t - (Y_{t-1}Y_{t-1}'\otimes I_K)\alpha] = 0 \tag{11}$$

and

$$\mathbb{E}(\Sigma_u - u_t u_t') = 0. \tag{12}$$

The empirical moments corresponding to the moment conditions for the structural parameters (8) and (10) are

$$\bar{m}^{\eta}(\eta,\gamma) = \frac{1}{T}\sum_{t=1}^{T} m_t^{\eta}(\eta,\gamma) \qquad \left(KK_1 + \tfrac{1}{2}K_1(K_1-1)\right)\times 1,$$

where $\eta = (\delta', \beta')'$ with

$$m_t^{\eta}(\eta,\gamma) = \begin{bmatrix} \mathrm{vec}\left(u_t(\alpha)z_t' - B_1\Sigma_{\mathbf{w}_1 z}\right)\\ \mathrm{vh}\left(B_1'\Sigma_u^{-1}u_t(\alpha)u_t(\alpha)'\Sigma_u^{-1}B_1\right)\end{bmatrix} \tag{13}$$

and for the reduced-form parameters we get empirical moments

$$\bar{m}^{\gamma}(\gamma) = \frac{1}{T}\sum_{t=1}^{T} m_t^{\gamma}(\gamma) \qquad \left(K(Kp+1) + \tfrac{1}{2}K(K+1)\right)\times 1$$

with

$$m_t^{\gamma}(\gamma) = \begin{bmatrix} (Y_{t-1}\otimes I_K)y_t - (Y_{t-1}Y_{t-1}'\otimes I_K)\alpha\\ \mathrm{vech}(\Sigma_u - u_t(\alpha)u_t(\alpha)')\end{bmatrix}. \tag{14}$$

Using a result by Crepon, Kramarz, and Trognon (1997, Proposition 1), we can set up an efficient GMM procedure by specifying the GMM objective function as

$$J(\eta) = T\bar{m}^{\eta}(\eta, \hat{\gamma})'\Omega(\hat{\eta}, \hat{\gamma})^{-1}\bar{m}^{\eta}(\eta, \hat{\gamma}), \tag{15}$$

where $\Omega(\eta, \gamma)$ is a suitable GMM weighting matrix, $\hat{\eta}$ is a consistent first-stage estimator of η and $\hat{\gamma}$ is a consistent estimator of γ that satisfies the condition

$$\bar{m}^{\gamma}(\hat{\gamma}) = 0. \tag{16}$$

In the present setup, we can use the LS estimator $\hat{\gamma}^{LS}$ as estimator for γ because it satisfies condition (16).

An asymptotically efficient estimator for η, denoted by $\hat{\eta}^{GMM}$ in the following, is obtained if the weighting matrix $\Omega(\eta, \gamma)$ is chosen as

$$\Omega(\hat{\eta}, \hat{\gamma}^{LS}) = \frac{1}{T}\sum_{t=1}^{T} \omega_t(\hat{\eta}, \hat{\gamma}^{LS})\omega_t(\hat{\eta}, \hat{\gamma}^{LS})', \tag{17}$$

where $\hat{\eta}$ is some consistent first-stage estimator of η and

$$\omega_t(\eta, \gamma) = m_t^{\eta}(\eta, \gamma) - \left(\frac{1}{T}\sum_{t=1}^{T}\frac{\partial m_t^{\eta}(\eta, \gamma)}{\partial \gamma'}\right) \left(\frac{1}{T}\sum_{t=1}^{T}\frac{\partial m_t^{\gamma}(\gamma)}{\partial \gamma'}\right)^{-1} m_t^{\gamma}(\gamma) \tag{18}$$

(see Crepon, Kramarz, and Trognon 1997). It is shown in Section OS.1 of the online supplement that

$$\omega_t(\eta, \hat{\gamma}^{LS}) = m_t^{\eta}(\eta, \hat{\gamma}^{LS}) - \begin{bmatrix} \left((\frac{1}{T}\sum_{t=1}^{T} z_t Y_{t-1}')(\frac{1}{T}\sum_{t=1}^{T} Y_{t-1}Y_{t-1}')^{-1} Y_{t-1} \otimes I_K\right)\hat{u}_t \\ -2\mathbf{S}_{K_1}\operatorname{vec}\left(B_1'\widehat{\Sigma}_u^{-1}(\widehat{\Sigma}_u - \hat{u}_t\hat{u}_t')\widehat{\Sigma}_u^{-1}B_1\right) \end{bmatrix} \tag{19}$$

is a closed-form of the correction term for our model if γ is replaced by the LS estimator $\hat{\gamma}^{LS}$. Here $\hat{u}_t$ denotes again reduced-form LS residuals and $\widehat{\Sigma}_u = T^{-1}\sum_{t=1}^{T}\hat{u}_t\hat{u}_t'$. For η we may choose a consistent first-stage estimator obtained, for example, by minimizing the GMM objective function (15) with $\Omega(\eta, \gamma)$ replaced by an identity matrix. The procedure can also be iterated, using the latest estimate of η and $\hat{\gamma}^{LS}$ in each step.

Of course, there is no need for determining the optimal weighting matrix in a just-identified case, where the GMM estimator is invariant to the weighting matrix. Such a case is, for instance, considered by Lakdawala (2019). Suppose there are $N = K_1$ proxies and we order the shocks such that the first proxy is correlated with the first shock, the second proxy is correlated with the second shock, etc.. Then, without loss of generality, we can scale the shocks such that the covariance matrix $\Sigma_{\mathbf{w}_1 z}$ has ones on its main diagonal. If it is triangular, it has then $\frac{1}{2}K_1(K_1 - 1)$ free elements and δ is a $\frac{1}{2}K_1(K_1 - 1)$-dimensional vector. As β is KK_1-dimensional, we have $\frac{1}{2}K_1(K_1 - 1) + KK_1$ parameters in η and this is precisely the number of moment conditions in (8) and (10). They just-identify η.

There can also be further zero restrictions on $\Sigma_{\mathbf{w}_1 z}$ if some of the proxies are uncorrelated with some of the $\mathbf{w}_{1t}$ shocks. As mentioned earlier, some authors assume that each proxy is correlated with one shock only (e.g., Altavilla et al. 2019). In other words, the proxies individually satisfy the relevance and exogeneity conditions (4)/(5). In that case, $\Sigma_{\mathbf{w}_1 z}$ is a $(K_1 \times K_1)$ identity matrix under our assumption that it has ones on its main diagonal. Thus, η consists of β only and the moment conditions are over-identifying. In the following we will primarily deal with that case.

Using this approach, we estimate the impact effects of the first K_1 structural shocks, scaled versions of which can be obtained as $\Sigma_{\mathbf{w}_1}^{-1}\mathbf{w}_{1t} = B_1'\Sigma_u^{-1}u_t$ from the reduced-form residuals. The GMM approach aims at estimating the structural shocks of interest in such a way that they are instantaneously uncorrelated. It should be clear that, if there are over-identifying moment conditions, the empirical correlation between the components of the estimated $\mathbf{w}_{1t}$ may be nonzero because $J(\hat{\eta}^{GMM}) > 0$. This feature can also be used to set up Hansen's J-test for misspecification using that, under standard GMM assumptions,

$$J(\hat{\eta}^{GMM}) \xrightarrow{d} \chi^2(\text{df}), \tag{20}$$

where df stands for the number of over-identifying moment conditions. If $\Sigma_{\mathbf{w}_1 z} = I_{K_1}$, df $= \frac{1}{2}K_1(K_1 - 1)$. Note that the asymptotic χ^2-distribution relies on the use of the optimal weighting matrix derived in Section OS.1 of the online supplement which corrects for the nuisance parameters. It assumes that the model and the moment conditions are correctly specified. Thus, it can be used as usual for checking for misspecification. The diagonality of $\Sigma_{\mathbf{w}_1 z}$ is one assumption that may not hold in practice and, diagnosing model misspecification by a Hansen test may be a signal of incorrect moment conditions. Of course, not rejecting the model specification by the J-test may be due to lack of power of the test. However, in the simulations in Section 3 we demonstrate that the J-test has good power in small samples against violations of the diagonality of $\Sigma_{\mathbf{w}_1 z}$. Moreover, our GMM approach enables the user to estimate $\Sigma_{\mathbf{w}_1 z}$ and also use other tools for assessing its diagonality, as we show in Section 4.

If the moment conditions are rejected, other, possibly less restrictive identification conditions may be considered. For example, one may use triangularity of $\Sigma_{\mathbf{w}_1 z}$ as in Lakdawala (2019) or exclusion restrictions on the impact effects as in Mertens and Ravn (2013). Of course, such just-identifying restrictions should only be used if there are good reasons for their validity, as they cannot be checked by statistical tests.

If we have standardized the shocks such that $\Sigma_{\mathbf{w}_1 z}$ is an identity matrix as in the previous discussion of the GMM approach, we can, of course, rescale the shocks to have the desired size for an empirical analysis. For example, if a monetary policy shock is identified that moves an interest rate on impact, we can rescale the column of $\hat{B}_1$ corresponding to the shock such that the interest rate changes by, say, 25 basis points on impact.

Our method is related to recent work by Hou (2024), who develops a Bayesian method which can accommodate the partial identification case we focus on here. In the frequentist paradigm, another approach for estimating the structural parameters B_1 is discussed by Angelini and Fanelli (2019). These authors assume a parametric model for the proxies and augment the VAR model by the proxies. Then they set up a minimum distance procedure that minimizes the distance of the structural parameters from the reduced-form parameters. Their approach also works for proxy VAR models where less than $K - 1$ shocks are identified by proxies. However, in addition to the B_1 matrix, the minimum distance method also estimates parameters of the model for the proxies. Given the way some proxies are constructed in the recent literature, their model is not a universally good approximation of the generating mechanism of the proxies. In particular, their model does not match the situation where observations for the proxy are only available at infrequent event dates and the proxy is set to zero on all other dates as, for example, in Piffer and Podstawski (2017), Wright (2012), Boer and Lütkepohl (2021), Gertler and Karadi (2015) and many other studies. Our GMM approach has the advantage of focussing on the parameters of interest, B_1, and it does not assume a specific model for the generating mechanism of the proxies and therefore also accommodates proxies with many zero values during the sample period.

3. Monte Carlo Study

3.1. Setup

An attractive feature of the orthogonality conditions introduced in the previous section is that they provide over-identifying moments, which are testable with the J-statistic. Unless the nuisance parameters are properly accounted for, however, this test will not follow the expected distribution. One might also expect the over-identified model to yield more precise estimates compared with the one-by-one proxy SVAR procedure since it incorporates additional information to estimate the same number of parameters. That both of these points are observed in finite samples is demonstrated with the following small Monte Carlo experiment.

The DGP is a VAR(1) process, $y_t = A_1 y_{t-1} + B\mathbf{w}_t$, with $\mathbf{w}_t \sim \mathcal{N}(0, \Sigma_\mathbf{w})$, where y_t and $\mathbf{w}_t$ are (3×1) vectors and $\Sigma_\mathbf{w}$ is a diagonal matrix. We fit VAR(4) models with a constant term and consider sample sizes, exclusive of pre-sample values, of $T = 100$ and 500 and the parameter matrices are

$$A_1 = \begin{bmatrix} 0.9 & 0 & 0 \\ 1/3 & 1/3 & 1/3 \\ 1/3 & 1/3 & 1/3 \end{bmatrix}, \quad B = \begin{bmatrix} 1 & 0.2 & 0.2 \\ 0.2 & 1 & 0.2 \\ 0.2 & 0.2 & 1 \end{bmatrix},$$
$$\Sigma_\mathbf{w} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \sigma_{w_3}^2 \end{bmatrix}.$$

The matrix A_1 has a maximum eigenvalue of 0.9 and, thus, the VAR process is stable but quite persistent. We set $\sigma_{w_3}^2 = 0.01$ or 1 and simulate 5000 Monte Carlo replications. The low value of $\sigma_{w_3}^2 = 0.01$ reflects an environment in which the two identified shocks account for a majority of the variation in all three variables. This case may be of practical interest as macroeconomists often study the most important sources of economic variation. When $\sigma_{w_3}^2 = 1$ all three shocks are equally important.

There are $K_1 = 2$ instruments available to identify the first two shocks. These two shocks are related to the vector of proxies, z_t, according to

$$z_t = \begin{bmatrix} 1 & \lambda \\ 0 & 1 \end{bmatrix} \mathbf{w}_{1t} + v_t, \quad v_t \sim \mathcal{N}(0, \Sigma_v), \tag{21}$$

where v_t is independent of $\mathbf{w}_{1t}$ and Σ_v is a diagonal matrix. We set the value of λ to zero initially such that the proxies are independent and each is correlated with one shock only. We will later vary λ to study the power of the J-test. The strength of the proxies as instrumental variables is determined by the correlations between the two identified shocks and their respective proxy variables, $\operatorname{corr}(w_{it}, z_{it}) = \sqrt{\operatorname{var}(w_{it})} \Big/ \sqrt{\operatorname{var}(w_{it}) + \operatorname{var}(v_{it})}$, for $\lambda = 0$. Since the two identified shocks have unit variance, we adjust the diagonal elements of Σ_v to achieve an intermediate correlation of 0.5 between the instruments and their associated shocks.

For each simulation we produce four estimates of the first two columns of B, that is for B_1. First, we apply the one-by-one proxy VAR procedure, which uses the two instrumental variables (proxies) but places no restriction on the correlation between the two identified shocks. Thus, it uses the moment conditions (8) such that the estimator for the impact effects is $\hat{B}_1(PVAR) = T^{-1}\sum_{t=1}^{T} \hat{u}_t z_t'$. Second, we apply the GMM method outlined in the previous section, which incorporates the over-identifying moment condition based on the orthogonality of the estimated shocks, and uses the adjusted GMM weighting matrix (17). The moment conditions for this estimator are (8) and (10). As first stage estimates for β and γ we use $\hat{\beta} = \operatorname{vec}(\hat{B}_1(PVAR))$ and $\hat{\gamma}^{LS}$, respectively. The resulting estimator will be referred to as the *adjusted GMM* estimator in the following. Third, we again use the moments (8) and (10) but use as GMM weighting matrix

$$\frac{1}{T}\sum_{t=1}^{T} m_t^{\beta}(\hat{\beta}, \hat{\gamma}^{LS}) m_t^{\beta}(\hat{\beta}, \hat{\gamma}^{LS})', \tag{22}$$

which would be the standard estimator of the weighting matrix if the nuisance parameters α and Σ_u are ignored. The corresponding estimator will be called the *unadjusted GMM* estimator. Finally, we compare to an estimator based on just-identifying moment conditions that assumes an upper-triangular $\Sigma_{\mathbf{w}_1 z}$ matrix and, hence, may be less efficient than our over-identified adjusted GMM. Its closed-form representation for the case where $\Sigma_{\mathbf{w}_1} = I_2$ is

$$\hat{B}_1 = \frac{1}{T}\sum_{t=1}^{T} \hat{u}_t z_t' \left[\operatorname{chol}\left(\left(\frac{1}{T}\sum_{t=1}^{T} z_t \hat{u}_t' \right) \widehat{\Sigma}_u^{-1} \left(\frac{1}{T}\sum_{t=1}^{T} \hat{u}_t z_t' \right) \right) \right]'^{-1} \tag{23}$$

(see Section OS.2 of the online supplement for a detailed derivation of this estimator). We refer to this as the *just-identified* estimator in the following.

3.2. Estimation Precision

Figure 1 shows kernel densities for the estimates of the six identified parameters of $B_1 = [b_{ij}]$ for each of the four estimators when we set $\sigma_{w_3}^2 = 0.01$. The vertical black lines correspond to the true values of the parameters. The following observations emerge from the figure:

1. The adjusted and unadjusted GMM estimators have very similar small sample densities. Therefore they cannot really be distinguished in the figure. In other words, for our limited

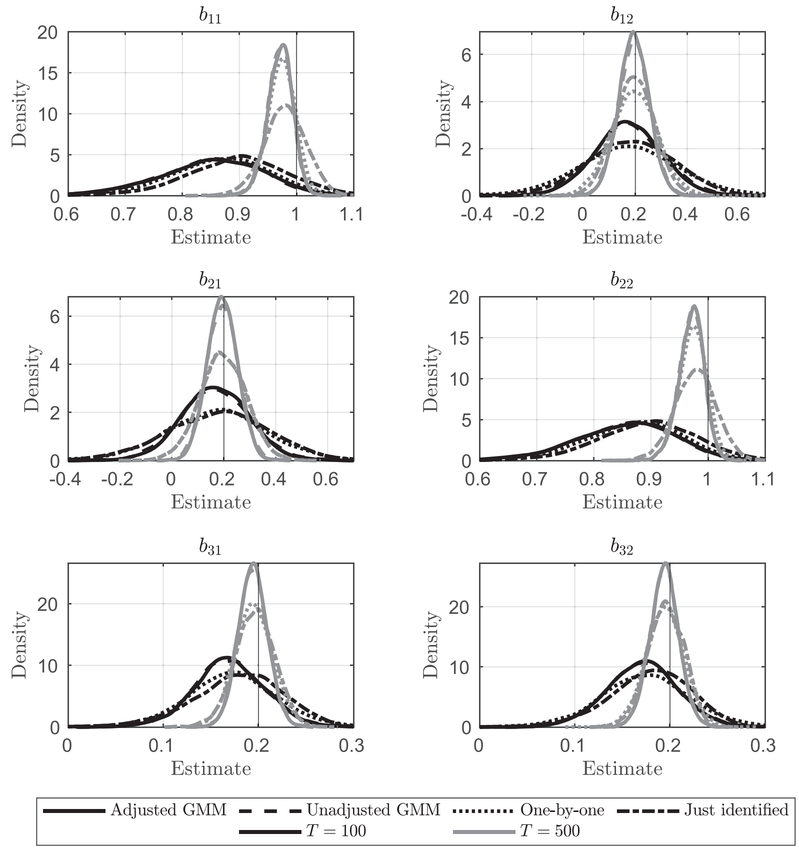

Figure 1. Monte Carlo kernel densities of estimates of $B_1 = [b_{ij}]$ when $\sigma^2_{w_3} = 0.01$.

simulation designs, the small sample properties of the estimators do not depend much on the adjustment of the GMM weighting matrix.
2. The two GMM estimators clearly dominate the one-by-one proxy VAR estimator and the just-identified estimator in that their densities are generally more or at least not less concentrated than the densities of the one-by-one proxy VAR estimates and, hence, the GMM estimators have smaller variances.
3. For some of the parameters the densities of all four estimators are concentrated around values different from the true parameter value in small samples. In other words, they are biased. The bias is similar for all four estimators and declines quickly with growing sample size.
4. There is no clear ranking of the one-by-one and just-identified estimators. For example, for b_{12} the just-identified estimator has a slightly more concentrated density for $T = 500$ than the one-by-one estimator while the situation is reversed for b_{22}.

The corresponding results for $\sigma^2_{w_3} = 1$ are depicted in Figure OS.1 in the online supplement and convey very similar results.

In summary, if each proxy is correlated with one shock only, the simulation results suggest that the over-identified GMM estimators should be used in applied work for estimating the structural parameters because they generally dominate the one-by-one and just-identified proxy VAR estimators in terms of small sample precision. In the examples in Section 4, we use the adjusted GMM procedure throughout because it has a clear advantage over the unadjusted GMM estimator when it comes to inference, as will be shown in Section 3.3.

Of course, in practice, researchers are not only interested in the impact effects of the shocks but perform, for example, impulse response analysis and, thus, the relative performance of our estimators in such an analysis is of interest. For comparing the one-by-one and adjusted GMM approaches in this context, we analyze coverage rates and confidence interval lengths for the impulse response functions. As the impact effects are part of the impulse responses at higher propagation horizons (see (2)), one might expect that also these impulse responses are estimated more efficiently with our GMM procedure.

Confidence intervals are calculated using the moving block bootstrap. That method was proposed by Jentsch and Lunsford (2019) and produces intervals that are robust to conditional heteroskedasticity. We use exactly the implementation described in Bruns and Lütkepohl (2023). It is also used in Section 4 where safeguarding against volatility changes in the residuals is desirable. For computational reasons, we consider only the adjusted GMM estimator and the one-by-one proxy SVAR estimator for the specification $\sigma^2_{w_3} = 0.01$ with sample sizes $T = 100$ and $T = 500$. We simulate 2000 Monte Carlo replications and 500 bootstrap replications.

It turns out that the coverage frequencies of both estimation methods are very similar (see Figure OS.2 in the online supplement) although the performance varies quite a bit for the different shocks, propagation horizons and sample sizes. The coverage rates are typically closer to the nominal 90% level for the larger sample size, as expected. Given that both estimation

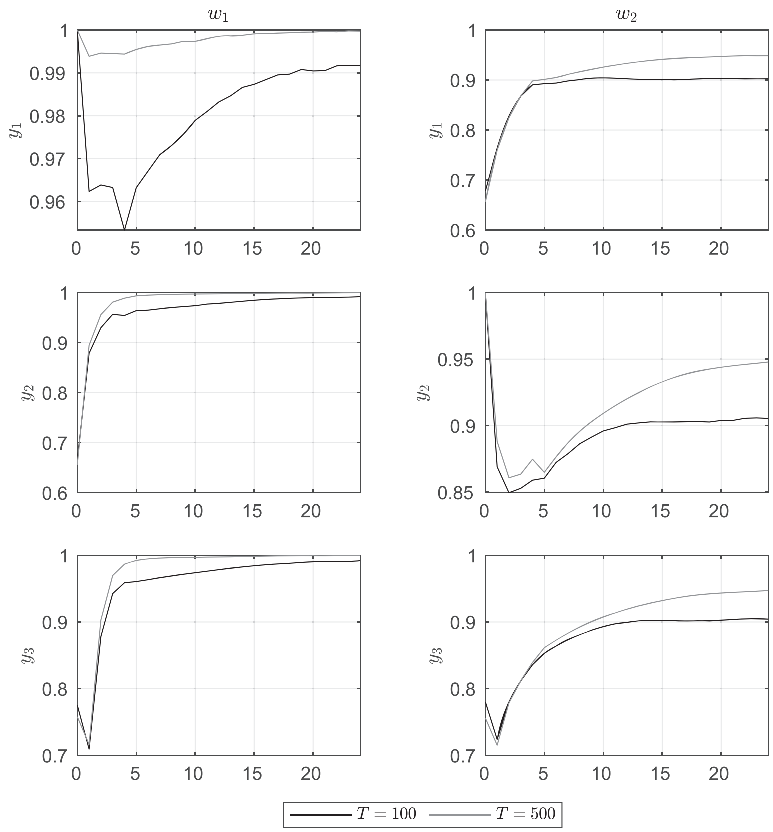

Figure 2. Monte Carlo relative average confidence interval lengths of the adjusted GMM over the one-by-one estimator.

methods achieve very similar coverage rates we can compare them on the basis of the confidence interval lengths.

Figure 2 shows ratios of the average bootstrap confidence interval lengths of the two estimators, where values less than one indicate narrower intervals for the adjusted GMM estimator. The GMM estimator achieves narrower confidence intervals at all horizons, as expected. The efficiency gains are largest at short horizons, in some cases as large as 30%, although in several cases these gains extend well out into the larger impulse response horizons. That is particularly so for the responses to the second shock, the effects of which are much less persistent than the first shock. At short impulse response horizons the efficiency gain for the GMM estimator is similar for the two sample sizes, indicating that our proposed estimator can provide more precise estimates even for larger sample sizes. Recall also that these results are obtained although in our simulation design there is only one over-identifying moment condition.

3.3. Small Sample Properties of the J-test

We will now explore the small sample properties of the J-test and its suitability for checking our assumed over-identifying moment conditions, based on the adjusted and unadjusted GMM estimators. Clearly, the J-test can be used only if there are over-identifying moment conditions and is not available for the other estimators. Table 1 shows rejection frequencies for the J-test calculated for the two GMM estimators at three popular significance levels. If the test statistic has the correct size, we would expect it to exceed the 10% critical value approximately 10% of the time, and likewise for any other critical value. We see that this is the case even for a relatively small sample size of $T = 100$ for the adjusted GMM but the test is severely undersized for the unadjusted GMM, indicating that a test of the over-identifying restrictions based on the latter estimator would under-reject. The rejection frequencies for the unadjusted GMM do not improve as the sample size increases and it is found for both values of $\sigma^2_{w_3}$. This illustrates that the adjustment for the GMM procedure outlined in the previous section is crucial for obtaining a reliable J-test.

To study the power of the J-test against incorrectly assuming diagonality of $\Sigma_{\mathbf{w}_1 z}$ we use a range of nonzero λ values in the DGP of z_t in (21) and show the corresponding rejection frequencies of the J-test for $T = 100$ in Table 2. While the adjusted estimator performs very well and has substantial small sample power even for samples as small as $T = 100$, the unadjusted estimator performs quite poorly. This is consistent with Table 1, which shows that the test statistic based on the unadjusted GMM is considerably undersized. The adjusted J-test has the advantage that it is performed conditionally on the reduced-form moment conditions to be satisfied and focusses on the over-identifying moment conditions for B_1. That may be one reason for its excellent small sample power against violations of the moment conditions related to B_1.

Overall our simulations show that if each proxy is correlated with one shock only, the adjusted GMM procedure should be used for estimating the structural parameters. Using the optimal GMM weighting matrix is crucial for taking full advantage of the inference possibilities that come with the GMM approach.

Table 1. Monte Carlo rejection frequencies of J-test under $\mathbb{H}_0$.

		$T=100$			$T=500$		
		10%	5%	1%	10%	5%	1%
$\sigma^2_{w_3}=0.01$	adjusted GMM	11.38%	5.72%	1.24%	11.22%	5.72%	1.32%
	unadjusted GMM	2.30%	0.68%	0.04%	1.84%	0.60%	0.02%
$\sigma^2_{w_3}=1$	adjusted GMM	11.38%	5.80%	1.18%	11.22%	5.70%	1.22%
	unadjusted GMM	2.24%	0.72%	0.06%	1.8%	0.60%	0.02%

Table 2. Monte Carlo rejection frequencies of the J-test at 10% level.

		λ					
		0	0.1	0.25	0.5	0.75	1
$\sigma^2_{w_3}=0.01$	adjusted GMM	12.86%	13.70%	24.96%	54.24%	80.46%	93.04%
	unadjusted GMM	2.60%	2.92%	7.76%	24.76%	50.02%	74.62%
$\sigma^2_{w_3}=1$	adjusted GMM	11.32%	14.16%	25.78%	55.66%	82.06%	93.74%
	unadjusted GMM	2.20%	2.86%	7.32%	24.80%	52.62%	76.44%

4. Empirical Examples

If a number of structural shocks are to be identified by multiple proxies and it can be defended that each proxy is correlated with a single shock only, then the adjusted GMM procedure offers a powerful tool for estimating the structural parameters and the corresponding J-test provides a way of checking the underlying identification assumptions. In the following we will illustrate by two examples how these tools work in practice.

4.1. U.S. Macroeconomic Shocks

Hou (2024) revisits the empirical analysis by Stock and Watson (2012) and investigates the response of U.S. macroeconomic variables to five shocks: An oil shock (w_t^{Oil}), a monetary policy shock (w_t^{MP}), a productivity shock (w_t^{Prod}), a liquidity and financial risk shock (w_t^{Risk}), and a fiscal policy shock (w_t^{Fiscal}). Hou (2024) employs nine proxies to identify these shocks in a 30-variate Bayesian SVAR model. As such a high-dimensional model would be a challenge for a frequentist analysis such as ours, we reduce the dimensionality and aim to identify the same five shocks using a seven-variate model with five proxies ($z_t^{Oil}, z_t^{MP}, z_t^{Prod}, z_t^{Risk}, z_t^{Fiscal}$). The variables we include represent standard macroeconomic variables that have been used in empirical macro models. Specifically, we include the oil price, the Fed Funds rate, output per hour of all persons, the S&P 500, real government consumption expenditures and gross investment, real GDP growth, and the CPI. Table OS.1 in the online supplement presents the variables and their transformations as well as the proxies. The sample runs from 1959q2 till 2011q2, implying $T = 209$ quarterly observations. Following Hou (2024) we employ four lags and an intercept term in the VAR model.

The empirical correlations of the proxies for our sample period are presented in Table OS.2 of the online supplement, where it can be seen that the proxies either have a statistically insignificant or a low correlation. Although this absent or low correlation among the proxies is no guarantee that each proxy is correlated with a single shock only, we follow Stock and Watson (2012) and use the proxies one-by-one to identify the five shocks. The correlations between the estimated shocks and the proxies and the shocks are presented in Table 3. The confidence intervals are generated by the bootstrap method presented in the online supplement, Section OS.3. Apparently, the resulting estimated shocks have a correlation of up to 0.633 for the case of $\hat{w}_t^{Fiscal}(PVAR)$ and $\hat{w}_t^{Prod}(PVAR)$ (see the upper panel of Table 3). In other words, they are partly highly correlated. The correlations between the proxies and the shocks in the lower panel of Table 3 do not provide a case for the assumption of a diagonal $\Sigma_{\mathbf{w}_1 z}$ matrix because, for example, z_t^{Prod} is significantly correlated with $\hat{w}_t^{Fiscal}(PVAR)$.

As, strictly speaking, the relevant $\Sigma_{\mathbf{w}_1 z}$ matrix to look at is the covariance matrix between the proxies and the orthogonal (uncorrelated) shocks, we have also determined shocks with our GMM procedure and present their correlations and correlations with the proxies in Table 4. (In our GMM algorithm, the weighting matrix $\Omega(\beta, \hat{\gamma}^{LS})$ of equation (17) used in the GMM objective function (15) is chosen iteratively, using as stopping rule a relative change of the objective function of less than 5%.) Not surprisingly, there is no significant correlation between the GMM shocks. Recall, however, that the GMM approach is based on the assumption of a diagonal $\Sigma_{\mathbf{w}_1 z}$ matrix which may not hold in this case. Therefore, it is useful to take a look at the estimated $\Sigma_{\mathbf{w}_1 z}$ matrix for the GMM shocks shown in the lower panel of Table 4. That matrix does not look exactly like a diagonal matrix since, for instance, the correlation between z_t^{Fiscal} and $\hat{w}_t^{Prod}(GMM)$ is significantly different from zero. However, we can use the J-test to formally test the diagonality assumption for $\Sigma_{\mathbf{w}_1 z}$. The J-statistic takes a value of 38.879 which corresponds to a p-value of less than 1% of the relevant $\chi^2(10)$ distribution. Hence, the J-test clearly rejects a proper specification.

These results indicate that additional assumptions are necessary for identifying the five shocks properly as uncorrelated shocks based on the present proxies. We are not trying to come up with such assumptions because it is our intention to illustrate how our GMM approach and the corresponding J-test work in practice. The next example also serves that purpose.

4.2. U.S. Monetary Policy

Jarociński and Karadi (2020) investigate the impact of monetary policy in the United States and the euro area. They consider two relevant shocks, a monetary policy shock which we denote by w_t^{mp} and a central bank information shock, denoted by w_t^{cbi} in the following. The first one captures conventional monetary pol-

Table 3. Empirical correlations of one-by-one proxy VAR shocks and proxies for the Hou (2024)/Stock and Watson (2012) example, sample period 1959q2 - 2011q2, with 95% bootstrap confidence intervals.

	$\hat{w}_t^{Oil}(PVAR)$	$\hat{w}_t^{MP}(PVAR)$	$\hat{w}_t^{Prod}(PVAR)$	$\hat{w}_t^{Risk}(PVAR)$	$\hat{w}_t^{Fiscal}(PVAR)$
$\hat{w}_t^{Oil}(PVAR)$	1	−0.586 (−0.763, −0.342)	−0.257 (−0.392, −0.010)	0.167 (0.032, 0.300)	−0.166 (−0.329, −0.001)
$\hat{w}_t^{MP}(PVAR)$		1	−0.051 (−0.210, 0.110)	−0.037 (−0.165, 0.111)	−0.034 (−0.160, 0.099)
$\hat{w}_t^{Prod}(PVAR)$			1	0.234 (0.089, 0.367)	0.633 (0.540, 0.711)
$\hat{w}_t^{Risk}(PVAR)$				1	0.134 (−0.005, 0.270)
$\hat{w}_t^{Fiscal}(PVAR)$					1
	$\hat{w}_t^{Oil}(PVAR)$	$\hat{w}_t^{MP}(PVAR)$	$\hat{w}_t^{Prod}(PVAR)$	$\hat{w}_t^{Risk}(PVAR)$	$\hat{w}_t^{Fiscal}(PVAR)$
z_t^{Oil}	0.222 (−0.049, 0.476)	−0.130 (−0.452, 0.211)	−0.057 (−0.158, 0.056)	0.037 (−0.119, 0.166)	−0.037 (−0.142, 0.060)
z_t^{MP}	−0.432 (−0.604, −0.238)	0.737 (0.601, 0.827)	−0.038 (−0.201, 0.121)	−0.028 (−0.173, 0.128)	−0.025 (−0.145, 0.099)
z_t^{Prod}	−0.190 (−0.345, −0.030)	−0.038 (−0.203, 0.140)	0.739 (0.656, 0.806)	0.173 (0.034, 0.307)	0.468 (0.369, 0.559)
z_t^{Risk}	0.047 (−0.073, 0.167)	−0.010 (−0.101, 0.084)	0.066 (−0.053, 0.181)	0.281 (0.167, 0.386)	0.038 (−0.099, 0.172)
z_t^{Fiscal}	−0.041 (−0.182, 0.121)	−0.008 (−0.098, 0.084)	0.155 (0.020, 0.291)	0.033 (−0.084, 0.164)	0.245 (0.102, 0.397)

Table 4. Empirical correlations of GMM shocks for the Hou (2024)/Stock and Watson (2012) example, sample period 1959q2 - 2011q2, with 95% bootstrap confidence intervals.

	$\hat{w}_t^{Oil}(GMM)$	$\hat{w}_t^{MP}(GMM)$	$\hat{w}_t^{Prod}(GMM)$	$\hat{w}_t^{Risk}(GMM)$	$\hat{w}_t^{Fiscal}(GMM)$
$\hat{w}_t^{Oil}(GMM)$	1	−0.059 (−0.286, 0.168)	−0.053 (−0.235, 0.136)	−0.013 (−0.141, 0.118)	0.019 (−0.147, 0.179)
$\hat{w}_t^{MP}(GMM)$		1	−0.007 (−0.169, 0.158)	−0.051 (−0.188, 0.105)	0.032 (−0.091, 0.159)
$\hat{w}_t^{Prod}(GMM)$			1	−0.010 (−0.145, 0.125)	−0.029 (−0.169, 0.112)
$\hat{w}_t^{Risk}(GMM)$				1	−0.107 (−0.251, 0.045)
$\hat{w}_t^{Fiscal}(GMM)$					1
	$\hat{w}_t^{Oil}(GMM)$	$\hat{w}_t^{MP}(GMM)$	$\hat{w}_t^{Prod}(GMM)$	$\hat{w}_t^{Risk}(GMM)$	$\hat{w}_t^{Fiscal}(GMM)$
z_t^{Oil}	0.158 (−0.037, 0.358)	−0.126 (−0.436, 0.202)	−0.056 (−0.157, 0.057)	0.032 (−0.128, 0.166)	−0.029 (−0.168, 0.099)
z_t^{MP}	−0.105 (−0.313, 0.102)	0.732 (0.596, 0.823)	−0.036 (−0.199, 0.122)	−0.051 (−0.203, 0.117)	0.015 (−0.092, 0.125)
z_t^{Prod}	−0.045 (−0.201, 0.116)	−0.006 (−0.173, 0.173)	0.738 (0.656, 0.806)	0.002 (−0.134, 0.145)	−0.026 (−0.154, 0.104)
z_t^{Risk}	0.044 (−0.085, 0.170)	0.001 (−0.093, 0.097)	0.063 (−0.057, 0.179)	0.265 (0.135, 0.384)	−0.021 (−0.183, 0.133)
z_t^{Fiscal}	0.039 (−0.113, 0.226)	0.007 (−0.082, 0.105)	0.156 (0.022, 0.292)	−0.015 (−0.128, 0.113)	0.175 (0.046, 0.328)

icy action such as changes in interest rates, while w_t^{cbi} captures the impact of the assessment of the economic outlook conveyed by the central bank. Jarociński and Karadi (2020) construct different sets of proxies z_t^{mp} and z_t^{cbi} to identify the shocks. Furthermore, they use sign restrictions to properly identify their shocks of interest and Bayesian methods to perform their analysis. We will focus on one of their U.S. models, a specific pair of proxies and frequentist methods, thereby deviating from Jarociński and Karadi (2020), to illustrate some of the points we have made in Section 2.

Our model involves five U.S. variables, the one-year constant-maturity Treasury yield, log S&P 500, log real GDP, the log GDP deflator, and the excess bond premium (EBP) as a measure for the recession risk in the next 12 months. We use monthly data from 1984m2 till 2016m12 from Jarociński and Karadi (2020), where further details on their construction are provided. The model fitted is a VAR(12) with intercept term.

The two proxies are constructed as follows: A series of Federal Funds futures surprises at the time of FOMC announcements is constructed and aggregated to monthly frequency. That monthly series is split up in two proxies by taking into account S&P 500 changes. When the S&P 500 moves in opposite direction to the Fed Funds futures, the Fed Funds futures surprise is classified as a value of z_t^{mp}, while it is classified as a value of z_t^{cbi} for all other periods. For all periods where no value is available, the proxies are set to zero. Thus, $z_t^{mp} z_t^{cbi} = 0$ by construction and,

Table 5. Empirical correlations of proxies and one-by-one proxy VAR shocks for Jarociński/Karadi (2020) example with 95% bootstrap confidence intervals.

	$\hat{w}_t^{cbi}(PVAR)$	$\hat{w}_t^{mp}(PVAR)$
z_t^{cbi}	0.203 (0.067, 0.299)	0.089 (−0.090, 0.242)
z_t^{mp}	0.101 (−0.018, 0.218)	0.232 (0.121, 0.335)
$\hat{w}_t^{cbi}(PVAR)$	1	0.436 (0.338, 0.534)
$\hat{w}_t^{mp}(PVAR)$		1

Table 6. Empirical correlations of proxies and GMM shocks for Jarociński/Karadi (2020) example with 95% bootstrap confidence intervals.

	$\hat{w}_t^{cbi}(GMM)$	$\hat{w}_t^{mp}(GMM)$
z_t^{cbi}	0.172 (−0.016, 0.338)	0.068 (−0, 115, 0.242)
z_t^{mp}	0.023 (−0.095, 0.138)	0.230 (0.123, 0.331)
$\hat{w}_t^{cbi}(GMM)$	1	−0.025 (−0.163, 0.124)
$\hat{w}_t^{mp}(GMM)$		1

as the proxies have nonzero means, their correlation is nonzero but small by construction.

In the upper panel of Table 5, we show the empirical correlations between the shocks and the proxies when the shocks are estimated using the one-by-one proxy VAR approach. The correlation between z_t^{cbi} and $\hat{w}_t^{mp}(PVAR)$ and between z_t^{mp} and $\hat{w}_t^{cbi}(PVAR)$ is small and not significantly different from zero. Given that the estimated $\Sigma_{\mathbf{w}_1 z}$ matrix is thus nearly diagonal, one may conclude that identifying the two shocks one-by-one may be justified. However, the resulting shocks have a significant correlation of 0.436 if the one-by-one proxy VAR approach is used for estimating them.

As it may be reasonable to assume in the present case that each of the proxies is only correlated with a single shock ($\Sigma_{\mathbf{w}_1 z}$ is diagonal), we can use our GMM procedure to obtain uncorrelated shocks. We have used an iterated weighting matrix $\Omega(\beta, \hat{\gamma}^{LS})$ in the GMM objective function (15) and obtained shocks similar to those from the one-by-one estimation. More precisely, $\hat{w}_t^{cbi}(PVAR)$ and $\hat{w}_t^{cbi}(GMM)$ are not quite as similar as $\hat{w}_t^{mp}(PVAR)$ and $\hat{w}_t^{mp}(GMM)$ (see Figure OS.3 in the online supplement).

The correlations of the GMM shocks and their correlations with the proxies are presented in Table 6. In this case, the empirical correlation between the estimated $\hat{w}_t^{cbi}(GMM)$ and $\hat{w}_t^{mp}(GMM)$ is -0.025 and, hence, very small and not significantly different from zero. The fact that the estimated correlation matrix corresponding to $\Sigma_{\mathbf{w}_1 z}$ based on the one-by-one proxy VAR shocks in Table 5 is diagonal is, of course, no insurance for getting also a diagonal $\Sigma_{\mathbf{w}_1 z}$ for the GMM shocks. Thus, we also present the estimated correlation matrix corresponding to $\Sigma_{\mathbf{w}_1 z}$ for these shocks in the upper panel of Table 6 and find that the assumption of a diagonal $\Sigma_{\mathbf{w}_1 z}$ matrix is supported by the very small and insignificant off-diagonal elements of the estimated correlation matrix. Thereby we also support the use of our GMM approach for estimation in this case. We have also performed the J-test and obtained a test value of $J = 1.346$ and a p-value of 0.25 of the corresponding $\chi^2(1)$ distribution which indicates that our test provides no evidence against the moment conditions being correct.

Given that an argument against using correlated shocks is that the corresponding impulse responses may reflect a distorted picture of the actual responses of the variables as isolated shocks are not likely to occur in practice, we present the impulse responses of the central bank information shock obtained with the one-by-one proxy VAR approach and the GMM approach in Figure 3. The shocks are scaled such that they increase the interest rate by 25 basis points on impact to make them comparable in size despite having potentially different variances. The confidence intervals around the impulse responses in Figure 3 are computed by the moving block bootstrap that was also used in the simulations in Section 3.2.

In Figure 3, the one-by-one proxy VAR impulse responses with 68% confidence intervals are shown on the left-hand side of the figure and the GMM confidence intervals are presented on the right-hand side. The point estimates obtained with both approaches are shown on the left-hand as well as the right-hand side to facilitate the comparison. In Figure 3, the confidence intervals of the GMM approach are overall somewhat smaller than the corresponding one-by-one proxy VAR intervals. Most point estimates of one approach are covered by the confidence intervals of the other approach. However, the responses of the stock index estimated by the one-by-one proxy VAR approach for the first 12 months after the shock do not fall inside the GMM confidence intervals and are much closer to zero than the responses estimated by GMM. In other words, for the orthogonalized central bank information shock, the impact effect is estimated to be much stronger on the stock market than indicated by the corresponding correlated shock. It is in fact quite plausible that the information released by the central bank is closely monitored by the stock market participants and, hence, the stock market response estimated by the GMM approach may be the more realistic one. Another striking difference is the response of EBP to the GMM shock compared to the shock estimated by the one-by-one approach. The confidence interval of the EBP response on impact to a GMM central bank information shock does not cover zero and, hence, one may conclude that the central bank can successfully reduce the risk of a recession by its communication despite increasing the interest rate. In contrast, relying on the one-by-one proxy VAR approach, zero is well inside the confidence intervals of propagation horizons of up to more than one year. Hence, in this case, one may underestimate the impact of the central bank communication shocks when using the correlated shocks.

The responses to $\hat{w}_t^{mp}(PVAR)$ and $\hat{w}_t^{mp}(GMM)$ shocks are presented in Figure OS.4 in the online supplement. They are also rather similar but display some noteworthy differences. For example, considering the point estimates, a 25 basis points interest rate shock is estimated by the GMM approach to have a stronger impact on the S&P 500 and the GDP deflator than the shock estimated by the one-by-one proxy VAR approach. Thus, considering uncorrelated shocks makes a difference for the impulse responses.

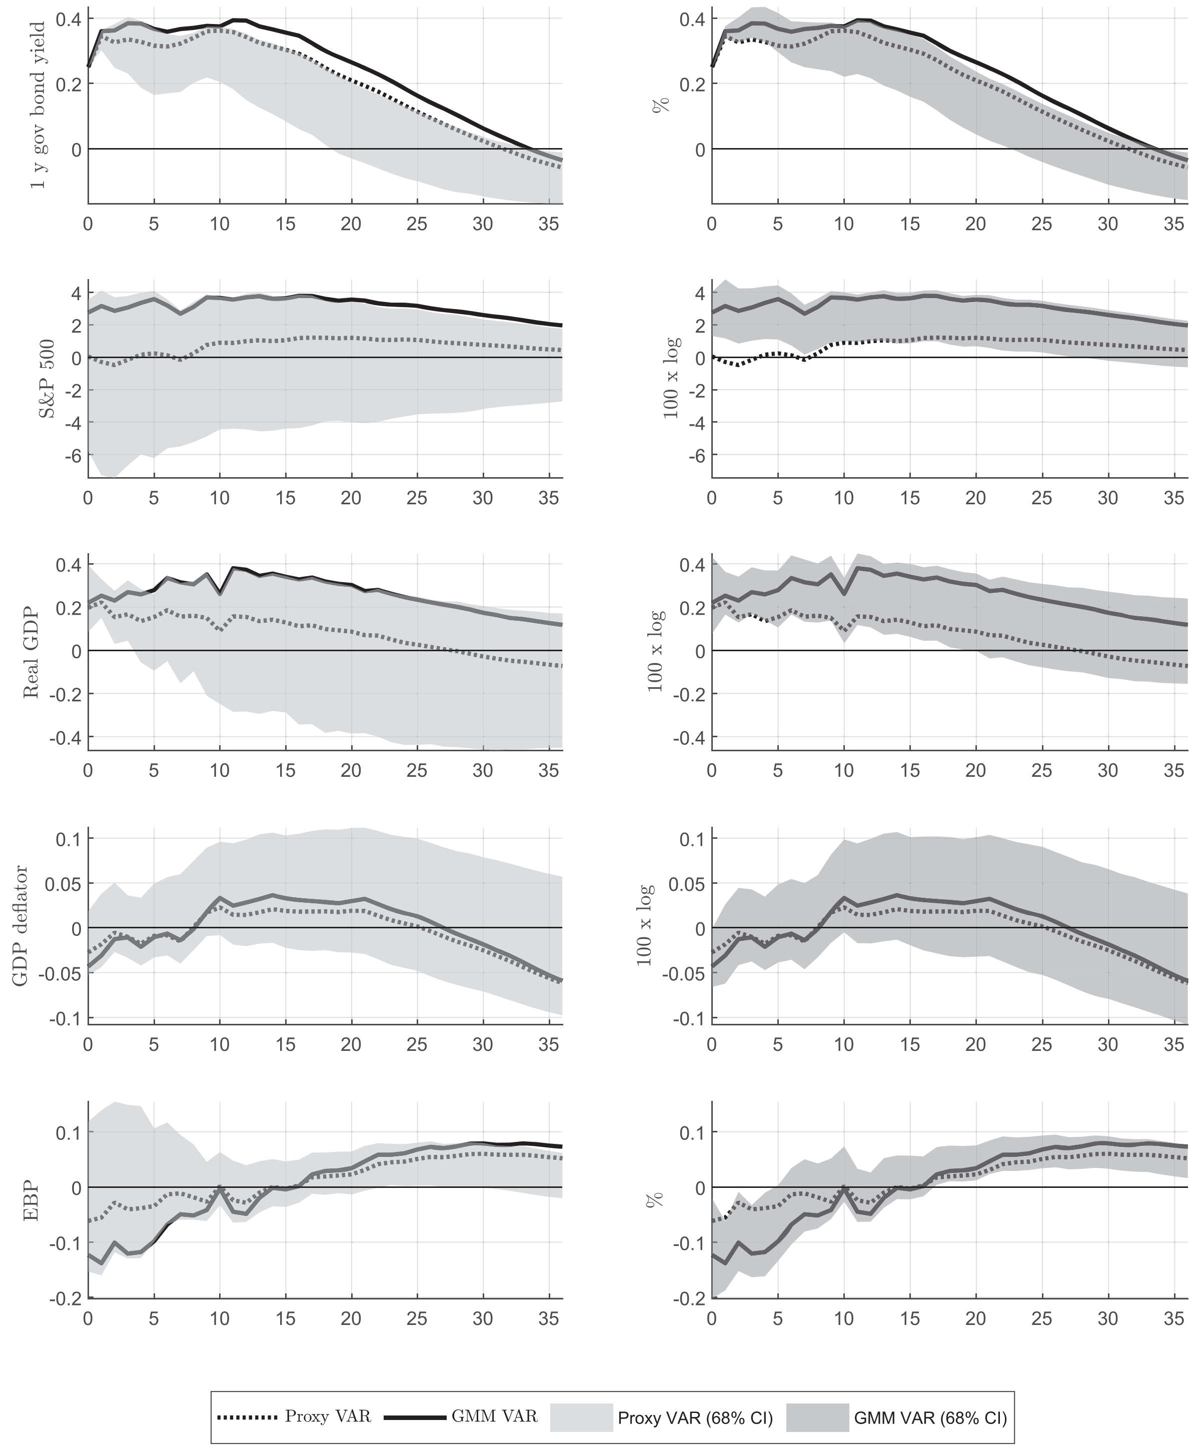

Figure 3. U.S. monetary policy example: comparison of impulse responses of a w_t^{cbi} shock estimated by the one-by-one proxy VAR approach (dotted lines, light grey confidence intervals) and with the GMM approach (solid lines, dark grey confidence intervals). The impulse responses are normalized to yield a 25 basis points increase in the one-year government bond yield on impact. The confidence intervals around the impulse responses are based on 5000 bootstrap samples.

This also holds for FEVDs. We present them for the PVAR and GMM shocks in Tables 7 and 8. In particular, the forecast error shares of the central bank information shock in Table 7 depend on the estimation method. For example, the share of forecast error variance of the log S&P 500 accounted for by $\hat{w}_t^{cbi}(PVAR)$ shocks is much smaller than that accounted for by $\hat{w}_t^{cbi}(GMM)$ shocks for all horizons shown in the table. A similar observation can be made for the forecast error shares of EBP. These results are, of course, well in line with the results for the impulse responses of these variables. Given the similarity of the $\hat{w}_t^{mp}(PVAR)$ and $\hat{w}_t^{mp}(GMM)$ in Figure OS.3, it is not surprising that the FEVDs of the two monetary policy shocks in Table 8 are rather similar.

It is perhaps worth drawing attention to one serious drawback of the correlated one-by-one proxy VAR shocks seen in Tables 7 and 8. The sum of the 1-step ahead forecast error variance shares

Table 7. Forecast error variance accounted for by Central Bank Information Shocks.

Shock	Variable	Forecast horizon						
		1	2	3	4	5	6	12
$\hat{w}_t^{cbi}(PVAR)$	1 y gov. bond yield	0.910	0.863	0.832	0.827	0.821	0.812	0.810
	log S&P 500	0.000	0.002	0.004	0.003	0.003	0.002	0.007
	log real GDP	0.083	0.144	0.147	0.150	0.142	0.135	0.099
	log GDP deflator	0.044	0.023	0.014	0.012	0.012	0.010	0.007
	EBP	0.039	0.045	0.037	0.036	0.033	0.031	0.027
$\hat{w}_t^{cbi}(GMM)$	1 y gov. bond yield	0.699	0.701	0.710	0.749	0.769	0.774	0.754
	log S&P 500	0.279	0.251	0.233	0.235	0.243	0.248	0.267
	log real GDP	0.079	0.141	0.167	0.196	0.211	0.219	0.242
	log GDP deflator	0.083	0.045	0.030	0.023	0.020	0.017	0.010
	EBP	0.120	0.172	0.164	0.173	0.169	0.166	0.141

Table 8. Forecast error variance accounted for by Monetary Policy Shocks.

Shock	Variable	Forecast horizon						
		1	2	3	4	5	6	12
$\hat{w}_t^{mp}(PVAR)$	1 y gov. bond yield	0.195	0.148	0.111	0.094	0.084	0.075	0.060
	log S&P 500	0.315	0.355	0.374	0.381	0.386	0.392	0.346
	log real GDP	0.075	0.065	0.106	0.132	0.168	0.196	0.293
	log GDP deflator	0.075	0.098	0.103	0.123	0.131	0.138	0.177
	EBP	0.430	0.455	0.461	0.456	0.446	0.434	0.401
$\hat{w}_t^{mp}(GMM)$	1 y gov. bond yield	0.114	0.079	0.054	0.042	0.035	0.030	0.020
	log S&P 500	0.399	0.439	0.457	0.466	0.475	0.483	0.436
	log real GDP	0.083	0.077	0.128	0.163	0.206	0.241	0.358
	log GDP deflator	0.056	0.083	0.091	0.112	0.120	0.128	0.177
	EBP	0.507	0.545	0.550	0.546	0.535	0.522	0.479

accounted for by $\hat{w}_t^{cbi}(PVAR)$ and $\hat{w}_t^{mp}(PVAR)$ (0.910 and 0.195, respectively) is greater than one which, by definition, is not possible for the actual shares of two shocks. Recall that the shares for all shocks should sum up to one. The reason is, of course, that the forecast error variance shares are computed conditional on the shocks being uncorrelated. Hence, performing an analysis with correlated shocks can lead to serious bias for the FEVDs. The problem does not occur when the shocks are estimated by our GMM approach (see Tables 7 and 8).

We emphasize again that Jarociński and Karadi (2020) use quite different identification and estimation methods. Therefore it is not surprising that their impulse responses differ from those obtained by our estimation approaches. We have deviated from their analysis to illustrate some of the theoretical points made in Section 2 of our paper.

5. Conclusions

This study shows that using proxies to identify structural shocks in a VAR analysis can lead to unintentionally correlated shocks. Such shocks are usually ruled out by assumption in structural VAR analysis because correlated shocks may lead to distorted impulse responses and FEVDs. When several proxies are used to identify a set of shocks and no further identifying information is available, in general, the proxies will identify only linear combinations of the impact effects. However, if the proxies satisfy the usual relevance and exogeneity conditions for valid proxies individually, in other words, if each proxy is correlated with exactly one shock only, using that feature and also imposing uncorrelatedness of the shocks implies even over-identifying restrictions for the shocks of interest. We have proposed a simple, efficient GMM approach that takes full advantage of the over-identifying restrictions and ensures uncorrelated shocks if each proxy is correlated with exactly one shock only. Our approach also provides a J-test with asymptotically valid χ^2 distribution. It also has good small sample properties and can be used to check the validity of the over-identifying moment conditions.

We present examples of structural VAR studies where multiple proxies identify more than one structural shock and where the structural shocks are not instantaneously uncorrelated if the proxies are used in the one-by-one way to identify the shocks. Enforcing uncorrelated shocks by using our GMM approach makes a difference for the impulse responses and FEVDs. Thereby we show that the problem of correlated shocks is relevant in practice and our GMM approach is a useful tool to overcome the problem.

Proxies have also been used to identify structural shocks in factor models (see, e.g., Stock and Watson 2012). Obviously, correlated shocks may also be obtained in that setting if the proxies are used one-by-one. Therefore, extending our analysis in that direction would be of interest. We leave it for future research because identifying structural shocks in factor models involves additional considerations (see, e.g., Kilian and Lütkepohl 2017, chap. 16).

Acknowledgments

The research presented in this article was partly circulated earlier in working papers entitled *Unintentionally Correlated Shocks in Proxy Vector Autoregressive Analysis* and *Over-identifying structural vector autoregressions with scarce external instruments*. We are grateful to Luca Fanelli, Allan W. Gregory, Gregor W. Smith, Jörg Breitung, Mikkel Plagborg-

Møller, James Hamilton, Kurt Lunsford, James Duffie, Nour Medahi, two anonymous referees, and the Editor Atsushi Inoue for helpful comments. Earlier versions of the research were presented at the annual conferences organized by the Royal Economics Society, April 11–13, 2022, the Canadian Economics Association, June 3–4, 2022, the International Association for Applied Econometrics, June 21–24, 2022, the Atlantic Canada Economics Association, October 14–16, 2022, the annual meeting of the Econometrics Committee of the Verein für Socialpolitik in Berlin, March 21–23, 2024, at the Workshop in Empirical Macroeconomics at King's College London, April 16, 2024, at the Annual Meeting of the Verein für Socialpolitik, Berlin, September 15–18, 2024, at the NBER-NSF Time Series Conference at University of Pennsylvania, September 20–21, 2024, The 11th Days of Econometrics for Finance: JEF'24, Rabat, Morocco, November 15–16, 2024, Macromodels International Conference 2024, Jablonna, Poland, November 19–21, 2024. We thank the participants for helpful comments.

Disclosure Statement

The authors report there are no competing interests to declare.

Funding

Martin Bruns thanks the British Academy for financial support, grant number TDA22\220046. James McNeil acknowledges the support of the Social Sciences and Humanities Research Council (SSHRC), Insight Development Grant number 430-2022-0007.

References

Altavilla, C., Brugnolini, L., Gürkaynak, R. S., Motto, R., and Ragusa, G. (2019), "Measuring Euro Area Monetary Policy," *Journal of Monetary Economics*, 108, 162–179. [1119,1122]

Angelini, G., and Fanelli, L. (2019), "Exogenous Uncertainty and the Identification of Structural Vector Autoregressions with External Instruments," *Journal of Applied Econometrics*, 34, 951–971. [1122]

Boer, L., and Lütkepohl, H. (2021), "Qualitative versus Quantitative External Information for Proxy Vector Autoregressive Analysis," *Journal of Economic Dynamics and Control*, 127, 104118. [1123]

Bruns, M., and Lütkepohl, H. (2022), "Comparison of Local Projection Estimators for Proxy Vector Autoregressions," *Journal of Economic Dynamics & Control*, 134, 104277. [1120]

——— (2023), "An Alternative Bootstrap for Proxy Vector Autoregressions," *Computational Economics*, 62, 1857–1882. [1124]

Crepon, B., Kramarz, F., and Trognon, A. (1997), "Parameters of Interest, Nuisance Parameters and Orthogonality Conditions: An Application to Autoregressive Error Component Models," *Journal of Econometrics*, 82, 135–156. [1119,1122]

Fanelli, L., and Marsi, A. (2022), "Sovereign Spreads and Unconventional Monetary Policy in the Euro Area: A Tale of Three Shocks," *European Economic Review*, 150, 104281. [1119]

Gertler, M., and Karadi, P. (2015), "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 7, 44–76. [1123]

Gregory, A. W., McNeil, J., and Smith, G. W. (2024), "US Fiscal Policy Shocks: Proxy-SVAR Overidentification via GMM," *Journal of Applied Econometrics*, 39, 607–619. [1119,1121]

Hou, C. (2024), "Large Bayesian SVARs with Linear Restrictions," *Journal of Econometrics*, 244, 105850. [1119,1120,1122,1126]

Jarociński, M., and Karadi, P. (2020), "Deconstructing Monetary Policy Surprises - The Role of Information Shocks," *American Economic Journal: Macroeconomics*, 12, 1–43. [1119,1121,1126,1127,1130]

Jentsch, C., and Lunsford, K. G. (2019), "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States: Comment," *American Economic Review*, 109, 2655–2678. [1124]

Känzig, D. R. (2021), "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements," *American Economic Review*, 111, 1092–1125. [1119]

Kilian, L., and Lütkepohl, H. (2017), *Structural Vector Autoregressive Analysis*, Cambridge: Cambridge University Press. [1130]

Lakdawala, A. (2019), "Decomposing the Effects of Monetary Policy Using an External Instruments SVAR," *Journal of Applied Econometrics*, 34, 934–950. [1119,1121,1122]

Lunsford, K. (2015), "Identifying Structural VARs with a Proxy Variable and a Test for a Weak Proxy," Technical Report, Federal Reserve Bank of Cleveland. [1119]

Lütkepohl, H. (2005), *New Introduction to Multiple Time Series Analysis*, Berlin: Springer-Verlag. [1120]

Mertens, K., and Ravn, M. O. (2013), "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States," *American Economic Review*, 103, 1212–1247. [1121,1122]

Piffer, M., and Podstawski, M. (2017), "Identifying Uncertainty Shocks Using the Price of Gold," *The Economic Journal*, 128, 3266–3284. [1119,1123]

Ramey, V. A. (2016), "Macroeconomic Shocks and their Propagation," in *Handbook of Macroeconomics* (Vol. 2), eds. J. B. Taylor and H. Uhlig, pp. 71–162, Amsterdam: Elsevier. [1119]

Stock, J. H., and Watson, M. W. (2012), "Disentangling the Channels of the 2007-09 Recession," *Brookings Papers on Economic Activity*, pp. 81–135. [1119,1126,1130]

——— (2018), "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments," *Economic Journal*, 128, 917–948. [1119,1120]

Wright, J. H. (2012), "What Does Monetary Policy Do to Long-Term Interest Rates at the Zero Lower Bound?" *Economic Journal*, 122, F447–F466. [1123]