

INSTITUTE  
OF ECONOMICS



Scuola Superiore  
Sant'Anna

LEM | Laboratory of Economics and Management

Institute of Economics  
Scuola Superiore Sant'Anna

Piazza Martiri della Libertà, 33 - 56127 Pisa, Italy  
ph. +39 050 88.33.43  
institute.economics@sssup.it

# LEM

## WORKING PAPER SERIES

### **Identification of Structural VAR Models via Independent Component Analysis: A Performance Evaluation Study**

Alessio Moneta <sup>a</sup>  
Gianluca Pallante <sup>a</sup>

<sup>a</sup> Institute of Economics and EMbeDS Department, Scuola Superiore Sant'Anna, Pisa, Italy.

**2020/24**

**September 2020**

**ISSN(ONLINE) 2284-0400**

# Identification of Structural VAR Models via Independent Component Analysis: A Performance Evaluation Study

Alessio MONETA\*<sup>1</sup> and Gianluca PALLANTE<sup>1</sup>

<sup>1</sup>Institute of Economics and EMbeDS  
Sant'Anna School of Advanced Studies  
Pisa, Italy

May 2, 2022

## Abstract

Independent Component Analysis (ICA) is a statistical method that linearly transforms a random vector. Under the assumption that the observed data are mixtures of non-Gaussian and independent processes, ICA is able to recover the underlying components, but at scale and order indeterminacy. Its application to structural vector autoregressive (SVAR) models allows the researcher to recover the impact of independent structural shocks on the observed series from estimated residuals. We analyze different ICA estimators, recently proposed within the field of SVAR analysis, and compare their performance in recovering structural coefficients. Moreover, we assess the size distortions of the estimators in hypothesis testing. We conduct our analysis by focusing on non-Gaussian distributional scenarios that get gradually close to the Gaussian case. The latter is the case where ICA methods fail to recover the independent components. Although the ICA estimators that we analyse show similar pattern of performance, two of them — the fastICA algorithm and the pseudo-maximum likelihood estimator — tend to perform relatively better in terms of variability, stability across sub- and super-Gaussian settings, and size distortion. We finally present an empirical illustration using US data to identify the effects of government spending and tax cuts on economic activity, thus providing an example where ICA techniques can be used for hypothesis testing.

*Keywords:* Independent Component Analysis, Identification, Structural VAR, Impulse response functions, Non-Gaussianity, Generalized normal distribution.

*JEL classification:* C14, C32, E62.

---

\*We thank Fulvio Corsi and Gabriele Fiorentini for comments and suggestions. We also acknowledge support by PRIN grant no. 20177FX2A7.

# 1 Introduction

The aim of this paper is to evaluate a set of methods that have been recently proposed to achieve statistical identification of structural autoregressive (SVAR) models based on non-Gaussianity. One of the most important and pursued objectives in macroeconomics is to estimate the dynamic effect of an unexpected change in one variable, usually called shock, on other variables. Since the seminal work of Sims (1980), the study of the joint dynamics of the main macroeconomic aggregates has been conducted in the framework of vector autoregressive (VAR) models. These models have been proposed as an alternative to large simultaneous equation models (Klein and Goldberger, 1955), which were criticized for their large number of identifying and arbitrary restrictions. However, while in forecasting (reduced-form) VAR models have been proven to be powerful tools, for policy analysis one needs to deal with a structural model.

It is indeed necessary to distinguish correlation from causation (Stock and Watson, 2017) if the goal of the analysis is measuring the effects of exogenous interventions on the system. Specifically, the residuals of an estimated (reduced-form) VAR model typically display cross-correlations, induced by contemporaneous causal relationships that cannot be detected in the regression estimates. There are infinite possibilities of linearly transforming the VAR model in order to get uncorrelated error terms, corresponding to infinite observational equivalent structural models. Researchers aim at finding the linear transformation that yields both uncorrelated (in some cases, independent) and economically meaningful disturbances, whose effects can be studied through impulse response analysis.

In the empirical macroeconomic literature, several identification criteria have been proposed. Following Sims (1980) and Sims (1986), many empirical works have exploited the Choleski decomposition of the covariance matrix of the VAR forecast errors. This procedure provides an orthogonalization of the residuals by imposing a recursive scheme on the contemporaneous causal structure, implicit in the ordering of the endogenous variables. The decision on the sequence of the variables is of crucial importance but sometimes loosely motivated. Economic theory or background knowledge may help achieve identification by imposing (typically zero) restrictions on the contemporaneous causal impact of one variable on another (Bernanke, 1986; Bernanke and Mihov, 1998; Blanchard and Perotti, 2002). The reliability of the implied causal structure, however, can

be hardly justified on the basis of mere *a priori* knowledge (Stock and Watson, 2001).

Alternative identification strategies are based on long-run restrictions, use of external instruments (extraneous data in general), sign restrictions, and heteroskedasticity (see Kilian and Lütkepohl, 2017; Stock and Watson, 2016, for an overview). External instruments have been used, for example, by Gertler and Karadi (2015), who identify the unexpected change in policy interest rate taking as instrument movements of futures prices around policy announcements. Romer and Romer (2010) adopt a similar identification procedure building a narrative series based on tax change announcements. Another popular identification approach is based on sign restrictions (see Uhlig, 2005; Mountford and Uhlig, 2009). A specific feature of this approach is that the structural coefficients are set identified, rather than point identified. It is also typical to rely on Bayesian methods of inference, which in sign-identified models may introduce the problem of priors that influence the posteriors of the structural coefficients (Kilian and Lütkepohl, 2017). Identification of SVAR models by heteroskedasticity is achieved by relying on the assumption that the contemporaneous causal structure does not vary over time, but their covariances change across regimes (Rigobon, 2003). The emerging of some of these identification strategies have made SVAR identification based less on restrictions guided by theory and more on statistical properties of the data.

In this respect, a recent stream of literature exploits directly specific statistical properties of the data — non-Gaussianity — for identification (see, among others, Lanne and Lütkepohl, 2010; Lanne et al., 2017; Gouriéroux et al., 2020; Lanne and Luoto, 2021; Fiorentini and Sentana, 2020; Guay, 2021). The idea of some of these studies is to recover the SVAR shocks as linear combinations of reduced-form VAR residuals disturbances, under the assumption that they are not just uncorrelated, but mutually statistical independent. This is possible by means of a purely data-driven set of statistical techniques, called Independent Component Analysis (ICA, see Comon, 1994; Hyvärinen, 1999; Eriksson and Koivunen, 2004). Under the assumptions of non-Gaussianity and independence of the shocks, the SVAR model is identified up to a re-scaling and re-ordering of the shocks. Empirical applications of SVAR analysis with ICA identification are flourishing (see, e.g., Moneta et al., 2013; Herwartz and Plödt, 2016; Capasso and Moneta, 2016; Gouriéroux et al., 2017; Herwartz, 2019; Maxand, 2020; Zema, 2021). Several different algorithms for learning independent components from data have been proposed and applied, espe-

cially in the fields of blind signal separation, neural networks, feature extraction (see, e.g., Hyvärinen and Oja, 2000), finance (see, e.g., Back and Weigend, 1997), causal inference and structural modeling (see, e.g., Shimizu et al., 2006).

In this paper we evaluate four methods of ICA. Notwithstanding the conspicuous number of ICA estimators emerged in the literature (see Acharya and Panda, 2008), we selectively focus on those that have been proposed or applied within the field of SVAR analysis, namely (1) the *fastICA* algorithm developed by Hyvärinen (1999) and employed by Moneta et al. (2013) and Guerini et al. (2020); (2) the minimization of the distance covariance proposed by Matteson and Tsay (2017); (3) the minimization of the Cramer-von-Mises statistics proposed by Herwartz and Plödt (2016); and (4) the pseudo-maximum likelihood estimator derived by Gouriéroux et al. (2017). Matteson and Tsay (2017), in fact, show how the distance-covariance method outperforms several ICA techniques under several distributional scenarios. Herwartz (2018) also undertakes a performance evaluation analysis, with a focus, however, on the discriminatory power of several identification schemes in detecting structural shocks embedded in a simple DSGE model for the Euro Area. Herwartz et al. (2021) compares different ICA methods with identification by heteroskedasticity. Our paper extends and completes the studies recently emerged in the literature by focusing on the statistical performance of the four methods mentioned above. Specifically, we study how these methods perform when the distribution of the structural disturbances gets gradually closer to the Gaussian case, which corresponds to the case in which the SVAR model cannot be identified through ICA.

We motivate the importance of analysing these settings with three considerations. Many macroeconomic time series, widely analyzed by means of VAR models, tend to display fat-tailed exponential distributions. However, their departure from normality is not always statistically significant as regards, for instance, OECD economies (Fagiolo et al., 2008). Secondly, it is often the case that different theoretical characterizations of firm dynamics and sectoral linkages move the distribution of economic activity away from the full-Gaussian case (Gabaix, 2011; Baqaee and Farhi, 2019). Finally, the log-transformation, widely implemented when estimating linearized multiplicative models or to stabilize the variance of a time series, is not enough to ensure a Gaussian approximation of the data generating process (Box and Cox, 1982; Nelson and Granger, 1979; Lütkepohl and Xu,

2012). We examine the different distributional scenarios through a  $p$ -generalized normal distribution. We are also able to analyze the distribution of the estimates of the parameters that are derived under the four different methods. Finally, we study the size distortions that arise when performing statistical inference on the coefficients of the impact multiplier matrix.

The paper is organized as follows: Section 2 presents the framework of our study, introducing the ICA-based SVAR model and the simulation exercise. Section 3 presents and discusses the results of our assessment. Section 4 discusses an empirical investigation in which the ICA-identified SVAR model is applied to study the effects of fiscal policy (government spending and tax cuts), using the data by Blanchard and Perotti (2002). Section 5 concludes.

## 2 The framework

### 2.1 SVAR and ICA

The SVAR model we study has the general form

$$A_0 y_t = c_t + \sum_{l=1}^q A_l y_{t-l} + \varepsilon_t, \quad (1)$$

in which  $q$  is the lag length,  $y$  is a  $k \times 1$  vector of endogenous variables,  $\varepsilon_t$  is a  $k \times 1$  vector of exogenous structural shocks  $(\varepsilon_{1t}, \dots, \varepsilon_{kt})'$ ,  $A_l$  is a  $k \times k$  matrix of parameters for  $0 \leq l \leq q$ ,  $c_t$  is a  $k \times 1$  vector of constants, which may also include a deterministic trend. (The model can also be easily extended to include exogenous variables.) We assume the  $\varepsilon_{1t}, \dots, \varepsilon_{kt}$  to be non-normally distributed (with at most one exception) and to be mutually independent, i.e.  $f(\varepsilon_{1t}, \dots, \varepsilon_{kt}) = f(\varepsilon_{1t}) \cdot \dots \cdot f(\varepsilon_{kt})$ , where  $f(\cdot)$  is the probability density function. We also assume  $A_0$  to be invertible. The model is structural because it is able to track the effect of statistically independent shocks on the endogenous variables, a crucial feature that makes the researcher able to identify, for example, the effect of a monetary or fiscal policy intervention.

The reduced-form representation implied by the structural model (1) is

$$y_t = d_t + \sum_{l=1}^q B_l y_{t-l} + u_t, \quad (2)$$

where  $B_l = A_0^{-1}A_l$  for  $1 \leq l \leq q$ ,  $d_t = A_0^{-1}c_t$ ,  $u_t = A_0^{-1}\varepsilon_t$ . Thus, we have that the reduced-form residuals  $u_t = (u_{1t}, \dots, u_{kt})$  are linear mixture of the structural shocks  $\varepsilon_t$ , namely:

$$u_t = B_0\varepsilon_t \iff \varepsilon_t = A_0u_t, \quad (3)$$

where  $B_0 = A_0^{-1}$ . Equation (3) is the model commonly studied in ICA, so that we refer to it as the ICA model. Using the ICA jargon, we call  $B_0$  the *mixing matrix*, since it linearly mixes the statistically independent components (i.e. shocks)  $\varepsilon_{1t}, \dots, \varepsilon_{kt}$ , and  $A_0$  the *unmixing matrix*.

Let us denote with  $a'_1, \dots, a'_k$  the rows of the matrix  $A_0$ . Any ICA procedure aims at estimating the  $k$ -length weight vectors  $a'_i$  for  $1 \leq i \leq k$ , which yield  $\varepsilon_{1t}, \dots, \varepsilon_{kt}$  as least dependent as possible. As proved by Comon (1994, Th. 11) and Eriksson and Koivunen (2004, Th. 3) (see also Gouriéroux et al., 2017), the independent components (shocks) are identifiable up to changes in scale (including sign) and ordering. More precisely, the matrix  $A_0$  in the ICA model (3) is identifiable up to the left multiplication by  $PD$ , where  $P$  is a permutation matrix and  $D$  a diagonal matrix with non-zero diagonal elements. Equivalently,  $B_0$  is identifiable up to the right multiplication by  $D^{-1}P'$  ( $P'$  is also a permutation matrix and  $D^{-1}$  a diagonal matrix).

ICA algorithms usually consist of two stages: a preliminary whitening and the actual ICA estimation. Whitening the data means to transform them so that they become uncorrelated and with unit variance. Suppose that we have estimated  $u_t$  and its non-diagonal covariance matrix  $\Sigma_u$ . Whitening can be obtained through the spectral (also called eigenvalue) decomposition or, as is popular in VAR analysis, via the Choleski factorization of  $\Sigma_u$ . The whitening transformation via the spectral decomposition consists of left multiplying  $u_t$  by  $(V\Lambda^{1/2})^{-1}$ , where  $V$  is the matrix containing the eigenvectors of  $\Sigma_u$ , and  $\Lambda$  is a diagonal matrix with the eigenvalues of  $\Sigma_u$  on the main diagonal. Whitening through the Choleski decomposition consists of left multiplying  $u_t$  by  $C$ , where  $C$  is the Cholesky factor of  $\Sigma_u$  (this can be done for any ordering of the variables). Without loss of generality, in this presentation of the ICA methods, we can directly assume that  $u_t$  is a vector of uncorrelated random variables (i.e.,  $u_t$  has already been whitened), so that the matrix  $B_0$  in equation (3) is orthogonal. Thus, the second stage of ICA estimation reduces to the problem of finding the rotation (orthogonal transformation) of the data  $u_t$  that delivers least dependent components  $\varepsilon_t$ .

We briefly review here four methods for ICA estimation, which we want to comparatively assess. Although further algorithms have been proposed in the literature (see, e.g., Cardoso, 1989; Hyvärinen, 2013), the approaches described below are good representative of the ICA methods discussed and applied in the econometric literature.

**1. FastICA.** A set of *fast* and *fixed-point* algorithms were proposed by Hyvärinen and Oja (1997, 2000) and Hyvärinen (1999). The *fastICA* approach is based on a fixed-point iteration scheme for finding a maximum of the non-Gaussianity of  $a'_i u_t$  (for  $i = 1, \dots, k$ ). It is called “fast” because it finds the maximally non-Gaussian components with a cubic convergence speed. As a measure of non-Gaussianity *fastICA* adopts an approximation of negentropy  $J(x)$ , a notion grounded on information theory. For a continuous random variable (or vector)  $x$  with density  $f(x)$ , negentropy is defined as  $J(x) = H(x_{gauss}) - H(x)$ , where  $x_{gauss}$  is a Gaussian random variable (or vector) of the same variance (covariance matrix) of  $x$ , and  $H(\cdot)$  is the differential entropy, i.e.  $H(x) = -\int f(x) \log f(x) dx$ . Such measure relies on the fact that a Gaussian random variable entails the largest entropy among all random variables of equal variance (Shannon, 1949). Hyvärinen and Oja (2000) also show that finding the most non-Gaussian directions  $a'_i u_t$  (for  $i = 1, \dots, k$ ) is equivalent to minimize the Kullback-Leibler divergence between the joint density  $f(a'_1 u_t, \dots, a'_k u_t)$  and the product of the marginal densities  $f(a'_1 u_t) \cdot \dots \cdot f(a'_k u_t)$ , which is a measure of mutual statistical dependence among the  $a'_i u_t$ 's and is also referred to as *mutual information*.

**2. Distance Covariance.** Matteson and Tsay (2017) propose to estimate the ICA model by finding a matrix of loadings  $A_0$  such that the distance covariance among the  $a'_i u_t$ 's is minimized. Distance covariance as measure of statistical dependence between random vectors was introduced by Székely et al. (2007).<sup>1</sup> Matteson and Tsay (2017) define an objective function to be minimized in function of  $\theta$ , which is the vector of angles defining a rotation matrix  $G(\theta)$ . Thus, the problem consists in finding  $\hat{\theta}$  such that the dependence (measured in terms of distance covariance) among the  $\varepsilon_{1t}, \dots, \varepsilon_{kt}$  that results from  $G(\hat{\theta})^{-1} u_t$  is minimized. Finally, the mixing matrix  $B_0$  is simply set to be equal to  $G(\hat{\theta})$ . In this approach, it is convenient to write  $G(\theta)$  as the product of  $k(k-1)/2$  distinct

<sup>1</sup>Let  $x^{(1)}$  and  $x^{(2)}$  be a  $d_1$ - and a  $d_2$ -dimensional random vectors. Let  $|\cdot|$  denote the Euclidean distance and let  $(\dot{x}^{(1)}, \dot{x}^{(2)})$  and  $(\ddot{x}^{(1)}, \ddot{x}^{(2)})$  be iid copies of  $(x^{(1)}, x^{(2)})$ . Székely et al. (2007) define the distance covariance between  $x^{(1)}$  and  $x^{(2)}$  as  $\mathcal{I}(x^{(1)}, x^{(2)}) = E|x^{(1)} - \dot{x}^{(1)}||x^{(2)} - \dot{x}^{(2)}| + E|x^{(1)} - \dot{x}^{(1)}|E|x^{(2)} - \dot{x}^{(2)}| - E|x^{(1)} - \dot{x}^{(1)}||x^{(2)} - \ddot{x}^{(2)}| - E|x^{(1)} - \ddot{x}^{(1)}||x^{(2)} - \dot{x}^{(2)}|$ .  $\mathcal{I}(x^{(1)}, x^{(2)}) = 0$  if and only if  $x^{(1)}$  and  $x^{(2)}$  are independent.



forms of Givens rotation matrices. In the 2-dimensional case we have only one angle to estimate:

$$G(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \quad (4)$$

In the three dimensional case we have 3 angles

$$G(\theta) = \begin{bmatrix} \cos \theta_1 & -\sin \theta_1 & 0 \\ \sin \theta_1 & \cos \theta_1 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \cos \theta_2 & 0 & -\sin \theta_2 \\ 0 & 1 & 0 \\ \sin \theta_2 & 0 & \cos \theta_2 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta_3 & -\sin \theta_3 \\ 0 & \sin \theta_3 & \cos \theta_3 \end{bmatrix} \quad (5)$$

For any  $k \times k$  matrix we have then  $k(k-1)/2$  rotation angles to estimate.

**3. Cramer-von-Mises distance.** Herwartz and Plödt (2016) and Herwartz (2018), similarly to Matteson and Tsay (2017), define an objective function to be minimized in function of  $\theta$  and exploits the same decomposition of  $G(\theta)$  in Givens matrices. But the minimization criterion is different. The selected vector of angles  $\hat{\theta}$ , which implies least dependent shocks, is the one that minimizes the value of the Cramer-von-Mises (CvM) statistics, developed by Genest et al. (2007). Specifically, this test statistics quantifies the distance between the empirical copula of the shocks vector  $\varepsilon_t = G(\theta)^{-1}u_t$  and the implied copula under mutual independence.

**4. Pseudo-maximum likelihood estimator.** This semi-parametric estimation method was proposed by Gouriéroux et al. (2017). It consists of a pseudo maximum likelihood (PML) estimator of the mixing matrix  $B_0$ , which maximizes the pseudo log-likelihood function, i.e.  $\mathcal{L}_T(B_0) = \sum_{t=1}^T \sum_{i=1}^k \log g_i(a'_i u_t)$ , where  $g_i(\cdot)$ 's are probability density functions, exploiting the condition that  $B_0$  is an orthogonal matrix. Gouriéroux et al. (2017) derive the asymptotic properties of the estimator, under possible specifications of  $\log g_i(\varepsilon_t)$ , but also assuming that some parts of the density functions may be misspecified.

## 2.2 Monte Carlo assessment

We study the performance of the just described ICA methods in estimating the model in equation (3) with  $k = \{2, 3\}$  and  $T = \{100, 200, 400\}$ , where  $k$  is the number of variables and  $T$  the sample size. The chosen set for  $k$  and  $T$  is due to the fact that we want to replicate a VAR model that is as close as possible to those commonly found in applied macroeconomics, where very long time series are seldom available to researchers. We

have in any case to restrict the analysis of the *CvM* method to relatively small sample sizes ( $T > 500$  is unfeasible, as Herwartz and Plödt (2016) point out), because the computational burden increases by an order of magnitude of  $O(T^2kn)$ , where  $n$  is the number of iterations implemented to generate the distribution of the statistics under the null of independence. We want to evaluate the performances of the ICA methods introduced above when the shocks/independent components in  $\varepsilon_t$ , and consequently (*a fortiori*) their linear combinations ( $u_t$ ), get gradually close to be normally distributed. In fact, it is often the case in empirical applications that the reduced-form residuals of an estimated VAR model turn out to be correlated and non-normal. But it may also be the case that normality of some of the residuals is not fully and clearly rejected, so that the researcher remains doubtful whether an ICA model can be legitimately applied for identification.

We perform the analysis by exploiting the properties of a class of exponential distributions, namely the  $p$ -generalized normal distribution (Box and Tiao, 1962; Goodman and Kotz, 1973): we let the underlying processes gradually approach to or diverge from a Gaussian distribution. In this manner we can analyze how the ICA procedures behave when the independent components diverge from normality both in the direction of a super-Gaussian (leptokurtic) and a sub-Gaussian (platykurtic) distribution. Often used for robustness studies (Subbotin, 1923; Box and Tiao, 1962; Tiao and Lund, 1970), this family of distributions has also been widely adopted in studies from different fields (e.g. signal processing, audio/video encoding, face recognition, finance), in which data often display non-Gaussian behavior (see Yu et al., 2012, for a review). Following the specification of Kalke and Richter (2013), a  $p$ -generalized normal distribution has a density function  $f$  of the form:

$$\text{g.norm} = f(x, p) = \frac{p^{1-1/p}}{2\Gamma(1/p)} \exp\left[-\frac{|x|^p}{p}\right] \quad x \in \mathbb{R}, \quad p > 0, \quad (6)$$

where  $\Gamma$  denotes the gamma function and  $p$  is a shape parameter that is informative about the rate of decay of the density function. With  $p = 2$ ,  $f(x, p)$  is a normal density function. Given this value, as  $p$  decreases, the distribution becomes more super-Gaussian, as  $p$  increases it becomes more sub-Gaussian. Specifically, with  $p = 0.5$ , equation (6) is the probability density function of a random variable with a Laplace distribution (a super-Gaussian distribution), with  $p = 100$  it corresponds to the case of a sub-Gaussian distribution, i.e. the uniform distribution. The two limiting cases,  $p = 0$  and  $p = +\infty$  correspond to a unit impulse function and to a real line, respectively.

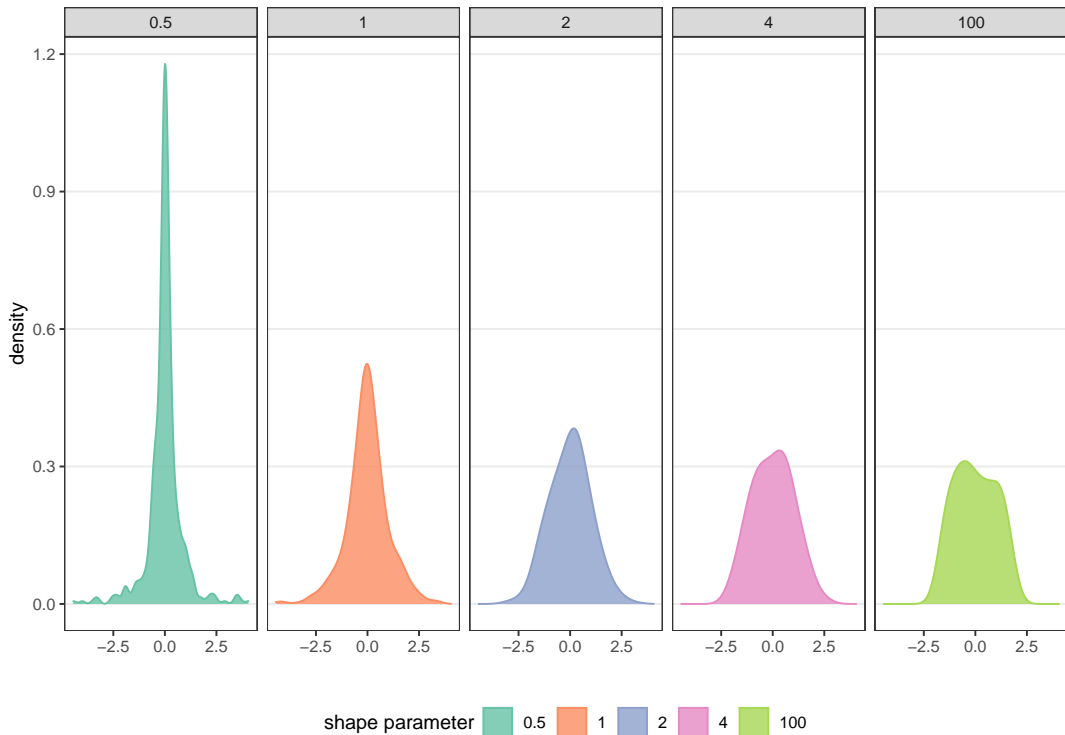


Figure 1: Kernel density estimates on data generated from the  $p$ -generalized normal distribution for different values of the shape parameter  $p$  (0.05, 1, 2, 4, 100): with  $p = 0.5$  super-Gaussian (Laplace),  $p = 2$  Gaussian,  $p = 100$  sub-Gaussian (uniform).

In our Monte Carlo experiment, we let the parameter  $p$  vary over a range of 20 values, 15 points uniformly located on the interval  $\{0.5, 3.5\}$  and 5 on  $\{4, 100\}$ : for each of these values we simulate  $k$  independent components.<sup>2</sup> As Figure 1 shows, for  $0.5 < p < 3.5$  the shape of the distribution changes substantially, while for  $p \geq 4$  the sub-Gaussian nature of the distribution is already pretty evident.

We split our Monte Carlo experiment into two different designs: a *general assessment* and a *specific assessment*. The *general assessment*, inspired by Matteson and Tsay (2017), aims at measuring the average performance of the four ICA methods in estimating the mixing matrix  $B_0$ , across random entries of the same matrix. Each Monte Carlo replication  $m$  generates a random ICA model  $u_t^{(m)} = B_0^{(m)} \varepsilon_t^{(m)}$ , where the components  $\varepsilon_t^{(m)}$  follow the  $p$ -generalized normal distribution (equation 6) with covariance matrix equal to the identity matrix, and  $B_0^{(m)}$  is a random  $k \times k$  mixing matrix with condition number  $1 \leq \mathcal{K}(B_0^{(m)}) \leq 2$ , simulated with the R-package ProDenICA (Hastie and Tibshirani, 2010).<sup>3</sup> Given the

<sup>2</sup>We use the R-package `rpgnorm` (Kalke, 2015).

<sup>3</sup>The condition number of a matrix  $B$ ,  $\mathcal{K}(B)$ , measures the “well-behavior” of  $B$ , namely the extent to which the solution  $x$  of the linear system  $Bx = c$  changes with the respect to changes in  $c$  (see e.g. Horn and Johnson, 2012, ch. 5.8) If  $\mathcal{K}(B) = 1$  the matrix  $B$  is said to be *perfectly conditioned*.

indeterminacy of the ICA model, namely its identification only up to column scale/sign and permutation of  $B_0$  (as mentioned in the previous sub-section), Matteson and Tsay (2017) suggest to measure the performance of the ICA methods with a metric, proposed by Ilmonen et al. (2010), that is invariant to this indeterminacy. Such measure is defined as follows:

$$D(B_0^{(m)}, \widehat{B_0^{(m)}}) = \frac{1}{\sqrt{k-1}} \inf \| C \widehat{B_0^{(m)}}^{-1} B_0^{(m)} - I_k \|_F \quad (7)$$

where  $B_0^{(m)}$  is the random matrix generated at replication  $m$ ,  $\widehat{B_0^{(m)}}$  is its estimate,  $C = P_{\pm} D_+$ , where  $P_{\pm}$  is any  $k \times k$  signed permutation matrix and  $D_+$  is any  $k \times k$  diagonal matrix with strictly positive diagonal element, and  $\|\cdot\|_F$  is the Frobenius norm.<sup>4</sup> The lower the index, the closer the estimate  $\widehat{B_0^{(m)}}$  to the true value  $B_0^{(m)}$ . We refer to this measure as the Minimum Distance Index (MDI).

The *specific assessment* aims at evaluating how well the ICA methods perform in identifying the structural impulse responses of a SVAR model, using different realizations of the same data generating process (dgp), for a given mixing matrix  $B_0$ . This also allows us to also compare, among each other, the Monte Carlo distributions of the parameter estimates derived by the four ICA methods and to analyze their statistical properties. The chosen mixing matrices for  $k = \{2, 3\}$  are, respectively:

$$B_0 = \begin{bmatrix} 1.14 & -0.38 \\ 0 & 1.26 \end{bmatrix}, \quad B_0 = \begin{bmatrix} 0.9 & 0.15 & 0.65 \\ -0.75 & 1.13 & 0.22 \\ 0.21 & -0.53 & 1.5 \end{bmatrix}. \quad (8)$$

Thus, in the case of two variables, we have an essentially triangular mixing matrix, while for  $k = 3$  we have a full matrix (all non-zero entries). This allows us to cover both a recursive and non-recursive mechanism of shocks' transmission. As mentioned in the previous sub-section, an ICA algorithm, which delivers a mixing (or unmixing) matrix, is not sufficient for full identification, since the mixing matrix is identified up to the right multiplication of  $DP$  ( $D$  is any diagonal matrix with all non-zero in the main diagonal and  $P$  is any permutation matrix).

While the matrix  $D$  can be appropriately chosen so that the each of the shocks/components  $\varepsilon_{it}$  in  $\varepsilon_t$  have variance one and impact positively on a specific variable  $y_{jt}$  in  $y_t$

---

<sup>4</sup>A signed permutation matrix is like a permutation matrix, with exactly one non-zero element for each row and column, but its non-zero elements are +1 or -1.

(for  $i$  and  $j$  in  $1, \dots, k$ ), there is not a general and unique method to choose  $P$ . Having estimated with an ICA algorithm the matrix  $\widehat{B}_0$ , there are different options:

1. One option is to apply the LiNGAM (Linear Non-Gaussian Acyclic Model) algorithm, as proposed by Shimizu et al. (2006), Hyvärinen et al. (2010) and Moneta et al. (2013). The permutation matrix  $P$  is found by assuming that the underlying contemporaneous causal structure is recursive (acyclic). This implies that  $\widehat{B}_0$  contains at least  $k(k-1)/2$  entries equal to 0 (but measurement errors) and that there exists only one  $P$  such that the main diagonal of  $B_0DP$  has all entries significantly different from 0.<sup>5</sup> The LiNGAM algorithm searches in the space spanned by all the possible permutation matrices the matrix  $\tilde{P}$  which minimizes a cost function that penalizes small absolute values in the main diagonal of  $\tilde{P}D\widehat{B}_0^{-1}$ . Furthermore, it searches for a permutation matrix  $\tilde{\tilde{P}}$  such that  $\tilde{\tilde{P}}\tilde{P}D\widehat{B}_0^{-1}\tilde{\tilde{P}}'$  is maximally close to a lower triangular matrix and sets its upper diagonal elements equal to zero.

2. The LiNG algorithm was formulated by Lacerda et al. (2008) (see Ciarli et al., 2019, for an application) as a variant of LiNGAM algorithm which relaxes the recursiveness (acyclicity) assumption, but still exploits the possibility that some entries of the unmixing matrix are equal to zero. Having tested (e.g., by bootstrap sampling) that some entries are vanishing, LiNG finds the set of permutation matrices  $C$ , such that for each  $\tilde{P} \in C$ ,  $\tilde{P}D\widehat{B}_0^{-1}$  does not contain zeros in the main diagonal. The permutation matrices satisfying this property, however, are, in general, not unique, so that the indeterminacy problem of ICA is not uniquely solved.

3. Having estimated  $\widehat{B}_0$ , and having re-scaled its columns by  $D$ , such that the shocks  $D^{-1}\widehat{B}_0^{-1}u_t$  have unit variance, another option is to apply what we call the *Maxfinder* criterion. This corresponds to find the permutation matrix  $\tilde{P}$  such that each column of  $\widehat{B}_0D\tilde{P}$  has off-diagonal elements that are smaller than the diagonal element. This criterion is guaranteed to give a (unique) result only under the assumption that the endogenous variable  $y_{jt}$  responds (within the sample period) more to  $\varepsilon_{jt}$  than to any  $\varepsilon_{it}$   $i \neq j$  in  $1, \dots, k$ . This fact would justify *labelling* the shock  $\varepsilon_{jt}$  as the  $y_{jt}$ -shock (e.g. the shock to tax revenue as the tax shock). In other terms, the *Maxfinder* criterion delivers a (unique) result only if each row of  $\widehat{B}_0D$  has exactly one maximum of any column of  $\widehat{B}_0D$ . If this is

---

<sup>5</sup>Equivalently, for any diagonal matrix  $D$ , there exists only one  $P$  such that the main diagonal of  $PD\widehat{B}_0^{-1}$  has all entries significantly different from 0.

not the case, for instance when a row contains the maximum entry of column 1 and 2, an option to solve this problem is to apply the *Maxfinder* criterion in a hierarchical manner,<sup>6</sup> as suggested by Bruns et al. (2021).

4. As mentioned in Berner et al. (2022), one can label the shocks identified by ICA using a criterion derived by the estimation of the forecast error variance decomposition. Following this approach, one finds a permutation matrix  $P$  such that the shocks  $\varepsilon_t = PD^{-1}\widehat{B}_0^{-1}u_t$  have the following characteristics: the prediction mean squared error (MSPE) of variable  $y_{1,t+h}$  is maximally accounted by shock  $\varepsilon_{1t}$ , the MSPE of  $y_{2,t+h}$  is maximally accounted by  $\varepsilon_{2t}$ , etc. This can be done for a specific time horizon  $h$  or summing up the contributions of the same shock across different time horizons (until a specific value of  $h$ ).

The list proposed above does not certainly exhaust the class of possible criteria for choosing the permutation matrix  $P$ .<sup>7</sup> It is also conceivable to apply a mixture of different criteria: for example one can apply the *MaxFinder* criterion and check whether it is consistent with the forecast error variance decomposition criterion (see Berner et al., 2022).

In our *specific assessment* we choose  $P$  according to the *MaxFinder* criterion, avoiding the recursiveness assumption and, at the same time, exploiting the fact that the mixing matrix in the underlying dgp (see equation 8) does not contain column maxima on the same row. This allows us to focus on the identification of the single coefficients. In the empirical application (Section 4), the assumption of column maxima on the same row is not guaranteed, so that in case of its failure, which is easily verifiable, we apply the step of LiNGAM (step 2 in the original algorithm by Shimizu et al., 2006), in which it is found a permutation matrix  $P$  that minimize the quantity  $1/\sum_i |w_{ii}|$  (penalizing small absolute values in the main diagonal), where  $w_{ii}$  is the  $(i, i)$  entry of the matrix  $W = PD\widehat{B}_0^{-1}$ .

---

<sup>6</sup>This hierarchical application of *Maxfinder* delivers a unique solution: (i) If the maximum entry of  $\widehat{B}_0D$  lies in position (row, column)  $(i, j)$ , then the  $j^{\text{th}}$  column becomes the  $i^{\text{th}}$  column in  $\widehat{B}_0DP$ , for a permutation matrix  $P$ . (ii) Repeat the same procedure starting from the output of the previous step,  $\widehat{B}_0DP$ , but compute the maximum entry neglecting all the entries lying in the  $i^{\text{th}}$  row and column of  $\widehat{B}_0DP$  until the  $k^{\text{th}}$  column has been permuted.

<sup>7</sup>Another criterion proposed in the literature is the sequence of transformations considered by Lanne et al. (2017), which imposes the entries of  $C = (\widehat{B}_0DP)$  to satisfy  $|c_{ii}| > |c_{ji}|$  for all  $i < j$ , where  $D$  is here a diagonal matrix that makes each column of  $\widehat{B}_0$  have Euclidean norm one.

## 3 Results

### 3.1 General assessment

Results about the performance of the four ICA estimators in the general assessment are shown in Figure 2. For each estimator, this figure displays the MDI, as specified in equation (7), across two different values of  $k$  ( $k = 2$  and  $k = 3$ ) and 20 values of  $p$ , fixing the sample size  $T = 400$ .<sup>8</sup> The thick lines in the plots trace the average MDI, calculated across 1000 distinct data generating processes, while the shadow areas display its (one standard deviation) variability. The performance of the four ICA methods is shown as the full Gaussianity ( $p = 2$ ) of the independent structural shocks is approached. As expected, when the  $\varepsilon_t$ 's are close to being normally distributed, all the methodologies score bad. To establish a negative benchmark, we calculated MDIs between a fixed  $B_0$  and 1000 random matrices of the same dimension (respecting the condition number as specified above). The 90% of these MDIs are above the dashed lines in each plot in Figure 2, which can therefore be considered as negative benchmark. For values of  $p$  close to Gaussianity, MDI averages are above (i.e. score worst) than this benchmark and the volatility of index increases.

On average, whatever the dimension of the system, *fastICA* and *PML* score relatively equal with the former slightly outperforming the latter when the  $\varepsilon$ 's are super-Gaussian. On the other hand, the opposite is true when the independent components (hereafter, ICs) are close from being normally distributed and sub-Gaussian. In such distributional scenario, *DCov* loses precision and score worse than any other method. *CvM* is dominated in almost all scenarios, performing better only when compared to *DCov* in the highly sub-Gaussian scenario ( $p > 4$ ). The results are compatible with and complementary to those in Matteson and Tsay (2017), where the ICs follow different families of distributions, both symmetric and asymmetric. Here instead we are interested in understanding the performances of ICA as we get closer to the full-Gaussian case. However, the exercise highlights that when the ICs are sub-Gaussian, *DCov* tends to have a lower performance.

---

<sup>8</sup>Results do not change qualitatively changing the sample size to  $T = \{100, 200\}$ . As expected though, the precision decreases as the sample size gets smaller (see Figure 9 in Appendix 5).

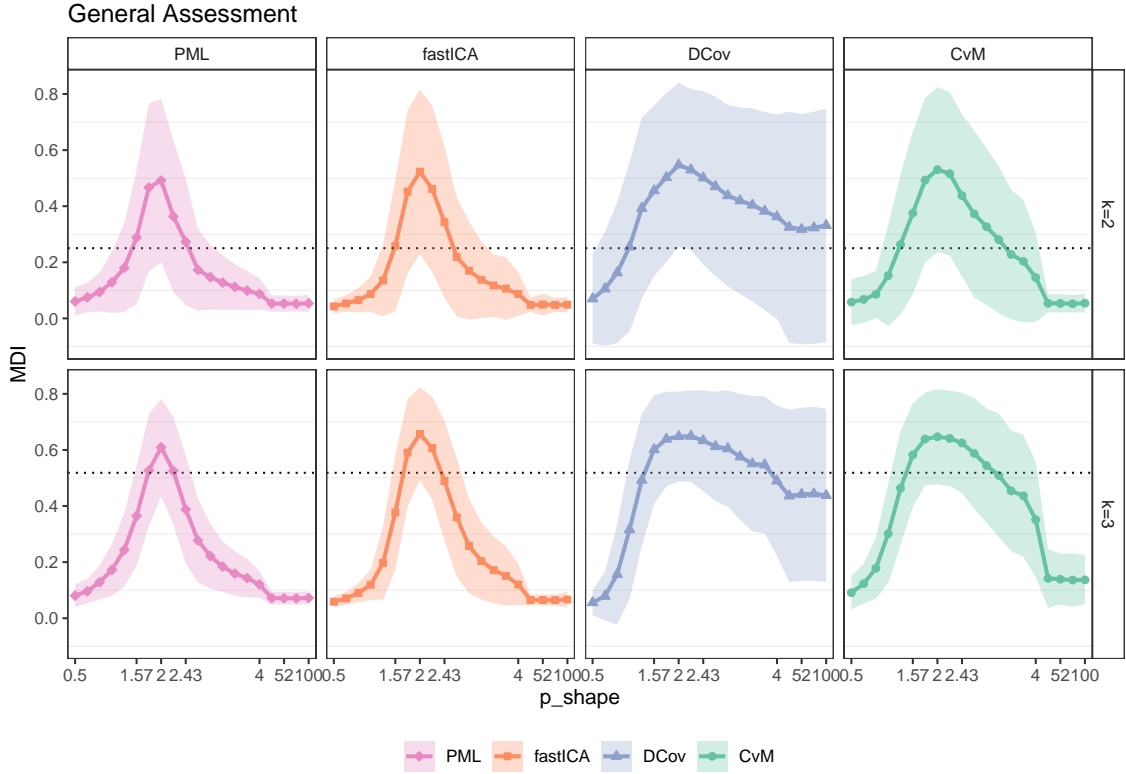


Figure 2: Each plot shows the MDI (thick lines with marks: averages across 1000 dgps), as specified in equation (7), for each of the four ICA estimators, across 20 values of  $p$ . The  $p$ -generalized normal distribution is super-Gaussian with  $p < 2$ , Gaussian with  $p = 2$ , sub-Gaussian with  $p > 2$ . Plots in the upper part of the figure correspond to the case  $k = 2$ , in the lower part to  $k = 3$ . Sample size  $T = 400$ . Shadow areas show one-standard deviation above and below the mean. Dashed lines constitute the negative benchmark: 90% of random estimates yield MDIs above this line.

### 3.2 Specific assessment

Having analysed the average performance of the four ICA methods, we turn now to study how well they perform when they are applied to recover the impact multiplier (mixing) matrix of a SVAR model. As mentioned in section 2.2, we focus on a specific data generating process. We artificially generate data  $u_t$  from the ICA model in equation (3) with shocks' covariance matrix  $\Sigma_\varepsilon = I$ ,  $c_t = 0_k$ , and  $A_0 = B_0^{-1}$ , where  $B_0$  is specified in equation (8) for  $k = 2, 3$ . Both specifications of  $B_0$  satisfy the condition that the maximum entry (in absolute value) of each column never appears on the same row. As mentioned in section 2.2, this assumption captures a feature that is often found in SVAR analysis, namely the fact that each structural shock tends to be mostly (contemporaneously) correlated with the variable in the system that is referring to and for which is labelled.



We apply the four ICA methods on the artificial  $u_t$ . We then apply the *MaxFinder* identification algorithm described in section 2.2, so to solve the indeterminacy that affects the ICA model. We then assess the performance of the ICA model in this specific exercise by using a variation of MDI: differently from equation (7), the index we use here measures the error between  $\widehat{B}_0$  and the known matrix  $B_0$  (this time fixed over Monte Carlo iterations, as specified in equation 8). Note that  $\widehat{B}_0$  is here column-permuted and scaled according the *MaxFinder* criterion. Hence, the distance index is now defined as

$$D(B_0, \widehat{B}_0) = \frac{1}{\sqrt{k-1}} \|\widehat{B}_0^{-1} B_0 - I_k\|_F \quad (9)$$

Figure 3 shows the MDI, as specified in equation (9), averaged across Monte Carlo replications, for different values of the parameter  $p$ , which determines the degree of “Gaussianity” of the ICs  $\varepsilon_t$ . To some extent, the results resemble those shown for the *General Assessment*. The performances of *fastICA* and *PML* are almost equivalent, with the difference that *fastICA* performs slightly better when the ICs are super-Gaussian and *PML* performs slightly better when the ICs are sub-Gaussian. One striking result is that the average value of the MDI for *DCov* is remarkably lower than for other estimators, but only for the two-dimensional case. Indeed, when  $k = 3$ , the behavior of the performance is very similar to the *general assessment*: MDI is highly volatile and its average does not improve significantly when the distribution of the ICs becomes sub-gaussian. The estimates of the mixing matrix under *CvM* display the largest variance among the estimators under analysis in most of the distributional scenarios that we have explored.

Although there are several measures of ICA performance proposed in the literature (see Nordhausen et al., 2011, for a review), those measures, like the one proposed in equations (7) and (9), are not informative about the distributional properties of the parameters’ estimates. As for any estimator, such properties are relevant when performing statistical inference: in our specific design, they are informative about the distribution of the entries of the mixing matrix  $B_0$ . Ultimately, these distributional properties may shed light on the contemporaneous causal relationships among the endogenous variables of the VAR system. Therefore, we study the Monte Carlo distribution of the errors between the entries  $\widehat{b}_{ij}$  of the estimated mixing matrix  $\widehat{B}_0$  and the entries  $b_{ij}$  of the known matrix  $B_0$  (for  $i, j = 1, \dots, k$ ).

Table 1 reports mean and standard deviation (across Monte Carlo runs) of  $\hat{b}_{ij} - b_{ij}$ , in four representative scenarios: (i) when the independent components  $\varepsilon_t$  have a strong super-Gaussian behavior ( $p = 0.5$ ); (ii) when Gaussianity is closer but the  $\varepsilon_t$  are still either super-Gaussian ( $p = 1.5$ ) or (iii) sub-Gaussian ( $p = 2.5$ ); (iv) when the  $\varepsilon_t$  follow an almost uniform distribution ( $p = 100$ ). We consider both  $k = 2$  (columns 1-8) and  $k = 3$  (columns 9-16). For  $k = 3$ , only the upper left block's parameters are considered. In almost all cases, the estimates are negatively biased (a result compatible with those of Gouriéroux et al., 2017). As expected, in those scenarios where full Gaussianity is closer, the bias is slightly more negative and the uncertainty of the estimates gets larger; the same holds when the dimension of the system ( $k = 3$ ) increases. For almost all parameters' estimates, in the extreme distributional scenarios ( $p = 0.5, 100$ ), all the methods score relatively equal; *fastICA* tends to have the smallest bias and, together with *DCov*, the smallest variance. In those cases close to the full-Gaussianity scenario, *DCov* performs relatively better when the dimension of the system is smaller, whilst *fastICA* seems to deliver better results when dimension of the system increases; *CvM* scores better in very few cases.

Finally, we evaluate the performance of the four ICA methods when statistical inference is conducted on the basis of a bootstrap procedure. Specifically, we compare bootstrap-based inference with the results derived from the pseudo-maximum likelihood approach. We expect that the *PML* estimator, derived by Gouriéroux et al. (2017) and drawn on asymptotic approximations, outperforms the bootstrap procedure, which we adopt since we do not know the asymptotic distribution of the estimate of  $B_0$ 's entries under *fastICA*, *DCov* and *CvM*. This should at least happen for large sample sizes. A small number of observations, however, may favour inference not based on asymptotic properties. Moreover, the comparison is interesting because it allows us to assess, from an another perspective, which method is more robust when the full-Gaussianity case is approached, in which, as stressed throughout the paper, ICA methods fail to recover the independent components.

The exercise is conducted in the following way. Again, we run a Monte Carlo exercise where we generate data from a ICA model with mixing matrix as specified in equation (8) and with shocks  $\varepsilon_t$   $p$ -generalized-normally distributed, for different values of  $p$  and sample size  $T = 400$ . For each Monte Carlo replication  $m = 1, \dots, 1000$  we estimate  $\hat{B}_0$  using one of the four ICA methods with the *MaxFinder* criterion.

For the *PML* method we derive  $\hat{\sigma}_{ij}$ , i.e the estimate of the standard deviation of  $\hat{b}_{ij}$  (for  $i, j = 1, \dots, k$ ), from the asymptotic covariance matrix of the mixing matrix found in Gouriéroux et al. (2017). On the basis of this, we build the  $(1 - \alpha)$ -confidence interval

$$C_m(\alpha) = [\hat{b}_{ij} - \phi_{\alpha/2}\hat{\sigma}_{ij}, \hat{b}_{ij} + \phi_{\alpha/2}\hat{\sigma}_{ij}], \quad (10)$$

with  $\phi_{\alpha/2}$  being the  $\alpha/2$ -quantile of the standard normal distribution. We do this for each Monte Carlo replication  $m$ , to which it corresponds a specific estimate  $\widehat{B}_0$ , with specific entries  $\hat{b}_{ij}$  (which we do not to index here with  $m$  just for simplicity) and corresponding confidence intervals  $C_m(\alpha)$ , and for each  $i, j$  in  $1, \dots, k$ . Finally, we calculate the frequency at which the true value  $b_{ij}$  (entries of the matrices specified in equation 8) falls in  $C_m(\alpha)$  across Monte Carlo replications. We should expect that the frequency of observing the true value  $b_{ij}$  in the confidence interval is equal to  $1 - \alpha$ .

As mentioned above, for the other three ICA methods we do not know the asymptotic distribution and therefore a bootstrap approach is necessary. Given the computational constraints of the exercise, we implement the *warp*-bootstrap, proposed by Giacomini et al. (2013), where it is shown that it is sufficient to have one bootstrap replication for each Monte Carlo run to obtain a reliable approximation of the statistics under analysis. The confidence interval is then built, for each  $i, j$  in  $1, \dots, k$ , in the following way:

$$C_m(\alpha) = [\hat{b}_{ij}^{(m)} - \hat{q}_{ij}(\alpha/2), \hat{b}_{ij}^{(m)} - \hat{q}_{ij}(1 - \alpha/2)], \quad (11)$$

in which  $\hat{q}_{ij}(\alpha)$  is the  $\alpha$ -quantile of the empirical distribution (across Monte Carlo runs) of  $\hat{b}_{ij}^{(m)*} - \hat{b}_{ij}^{(m)}$ , where  $\hat{b}_{ij}^{(m)*}$  is the estimate of the  $(i, j)$  entry of  $B_0$  obtained at the (unique) bootstrap draw corresponding to the Monte Carlo run  $m$ , and  $\hat{b}_{ij}^{(m)}$  is the estimate of the  $(i, j)$  entry of the mixing matrix  $B_0$  at the Monte Carlo run  $m$ . We then compute the frequency at which  $C_m(\alpha)$  contains the true value  $b_{ij}$  across the Monte Carlo replications. If the bootstrap procedure is consistent, we should then expect that the frequency at which the bootstrap-based confidence interval contains the true value (i.e. its the empirical coverage  $1 - \hat{\alpha}$ ) to be exactly equal to  $1 - \alpha$ , the nominal coverage.

Pooling all the entries of the mixing matrices' estimates of Equation 8, Figure 4 shows the distribution of the size distortions, measured by the ratio between the empirical coverage  $1 - \hat{\alpha}$  and its nominal counterpart  $1 - \alpha$ , of each estimator for values of  $(1 - \alpha) = \{0.99, 0.95, 0.90, 0.75, 0.50\}$ . The first result is that bootstrap-based confidence

intervals built under *fastICA* display the smallest size distortion, notwithstanding the dimension  $k$  of the system. Surprisingly, this is true even when compared with *PML*-based inference, especially in the case in which the independent components follow sub-Gaussian distributions ( $p = \{2.43, 100\}$ ). However, as stated in Gouriéroux et al. (2017), the choice of the pseudo-likelihood does indeed matter for the asymptotic accuracy of the *PML* estimator. Furthermore, all estimators but *CvM*, which also display the worst performance, tend to contain much more often than expected the true value of the parameter ( $\hat{\alpha} < \alpha$ ): this distortion decreases as the nominal size  $\alpha$  increases.

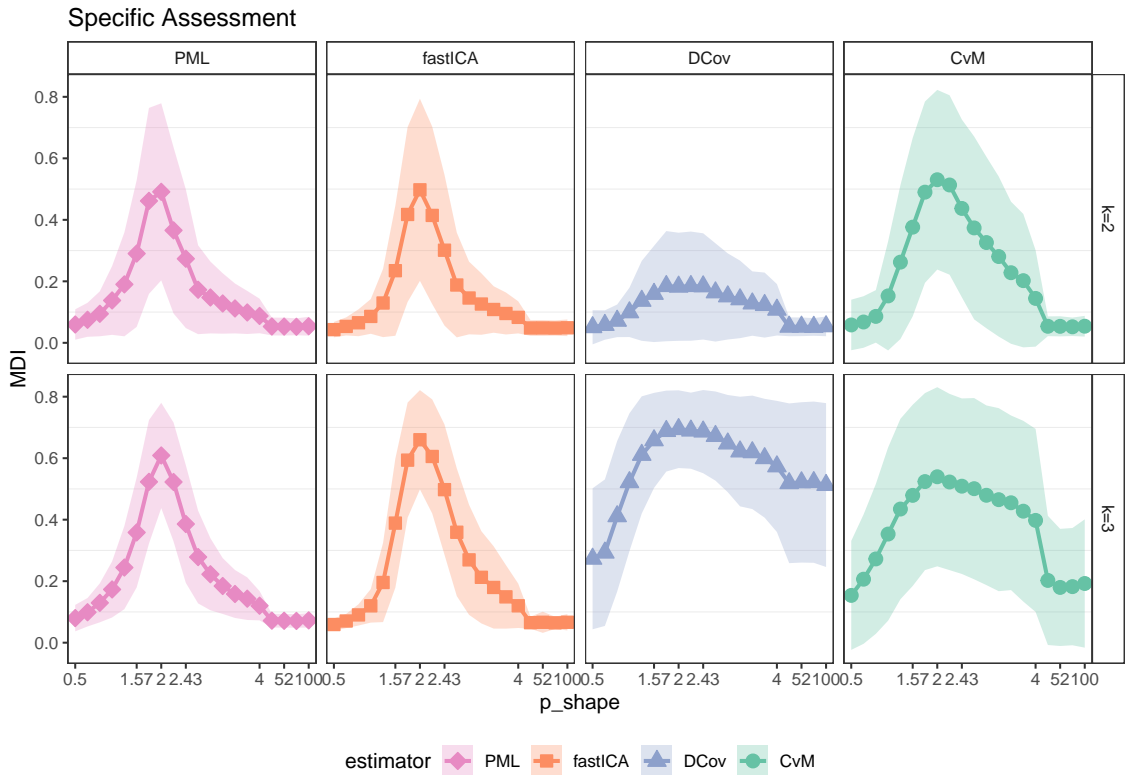


Figure 3: Each plot shows the MDI (thick lines with marks: averages across 1000 Monte Carlo replications), as specified in equation (9), for each of the four ICA estimators, across 20 values of  $p$ . The mixing matrix in the *dgp* is fixed as in equation (8). The  $p$ -generalized normal distribution is super-Gaussian with  $p < 2$ , Gaussian with  $p = 2$ , sub-Gaussian with  $p > 2$ . Plots in the upper part of the figure correspond to the case  $k = 2$ , in the lower part to  $k = 3$ . Sample size  $T = 400$ . Shadow areas show one-standard deviation above and below the mean.

## 4 Empirical application

In this section we discuss a macroeconomic application of the ICA approach to SVAR analysis with the aim of showing its potentials and challenges, having taken into account the

Table 1: Summary statistics of the errors between entries of  $\widehat{B}_0$  and of  $B_0$  (the latter as specified in equation 8). Sample size  $T = 400$ .

		$k = 2$								$k = 3$							
		$p = 0.5$		$p = 1.57$		$p = 2.47$		$p = 100$		$p = 0.5$		$p = 1.57$		$p = 2.47$		$p = 100$	
		mean	sd	mean	sd	mean	sd	mean	sd	mean	sd	mean	sd	mean	sd	mean	sd
<i>PML</i>	$\widehat{b}_{11} - b_{11}$	-0.002	0.007	-0.045	0.077	-0.039	0.068	-0.001	0.002	-0.005	0.042	-0.032	0.145	-0.025	0.147	-0.004	0.032
<i>fastICA</i>		-0.001	0.001	-0.032	0.067	-0.049	0.081	-0.001	0.001	-0.002	0.027	-0.020	0.144	-0.007	0.165	-0.003	0.031
<i>DCov</i>		-0.001	0.006	-0.015	0.034	-0.018	0.037	-0.001	0.002	0.032	0.138	0.116	0.160	0.135	0.149	0.066	0.146
<i>CvM</i>		-0.003	0.018	-0.076	0.106	-0.093	0.112	-0.001	0.002	-0.001	0.044	-0.014	0.177	-0.005	0.177	-0.007	0.064
<i>PML</i>	$\widehat{b}_{12} - b_{12}$	-0.002	0.065	-0.029	0.308	-0.035	0.287	0.001	0.051	-0.002	0.063	-0.067	0.310	-0.102	0.330	0.000	0.052
<i>fastICA</i>		0.000	0.038	-0.016	0.261	-0.044	0.317	0.001	0.046	0.000	0.041	-0.087	0.329	-0.144	0.384	0.000	0.047
<i>DCov</i>		0.001	0.051	-0.009	0.179	-0.017	0.200	0.001	0.052	-0.057	0.234	-0.179	0.306	-0.169	0.282	-0.127	0.314
<i>CvM</i>		0.001	0.079	-0.080	0.386	-0.112	0.421	0.001	0.052	-0.002	0.073	-0.186	0.433	-0.199	0.421	-0.013	0.108
<i>PML</i>	$\widehat{b}_{21} - b_{21}$	0.002	0.068	0.048	0.349	0.054	0.327	0.001	0.054	0.002	0.061	0.165	0.393	0.220	0.429	0.002	0.051
<i>fastICA</i>		-0.001	0.046	0.030	0.296	0.067	0.364	0.001	0.049	0.000	0.043	0.214	0.426	0.352	0.484	0.002	0.046
<i>DCov</i>		-0.002	0.077	0.015	0.205	0.028	0.231	0.002	0.054	0.164	0.329	0.458	0.371	0.458	0.344	0.388	0.361
<i>CvM</i>		-0.001	0.094	0.115	0.445	0.156	0.485	0.002	0.056	0.005	0.075	0.457	0.523	0.510	0.503	0.014	0.114
<i>PML</i>	$\widehat{b}_{22} - b_{22}$	-0.001	0.021	-0.041	0.110	-0.032	0.097	-0.002	0.017	-0.002	0.041	-0.025	0.142	-0.022	0.148	-0.003	0.035
<i>fastICA</i>		-0.001	0.013	-0.030	0.091	-0.040	0.106	-0.001	0.016	-0.002	0.026	-0.025	0.144	-0.017	0.157	-0.003	0.031
<i>DCov</i>		-0.003	0.022	-0.014	0.062	-0.016	0.074	-0.001	0.017	0.011	0.090	0.105	0.134	0.115	0.130	0.039	0.103
<i>CvM</i>		-0.004	0.033	-0.056	0.119	-0.065	0.131	-0.001	0.018	-0.003	0.045	-0.023	0.164	-0.011	0.174	-0.008	0.066

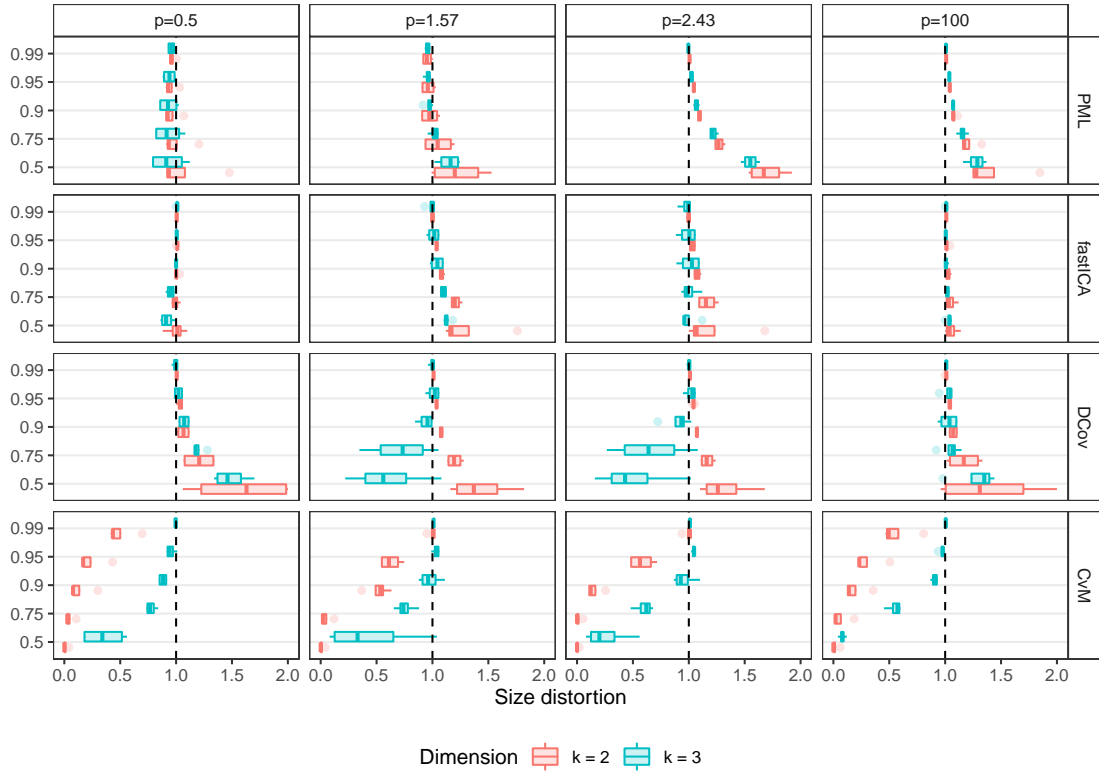


Figure 4: Distribution of size distortions (measured on the  $x$ -axis by the ratio between the empirical coverage  $1 - \widehat{\alpha}$  and its nominal counterpart  $1 - \alpha$ ) in drawing confidence intervals for the estimates of the mixing matrix entries, for different estimators (row-wise ordered panels), different distributional scenarios (column-wise ordered panels), when the dimension of the system  $k$  changes. Confidence intervals are constructed at nominal  $\alpha$  significance level. For *PML*, we use the asymptotic approximation derived by Gouriéroux et al. (2017), while for *fastICA*, *DCov* and *CvM* we implement the *warp*-bootstrap procedure described in text.

results of the previous section. ICA offers the opportunity to statistically test identifying restrictions (not limited to those that are over-identifying) on the coefficients of the impact (mixing) matrix. We consider here the very influential work on fiscal policy by Blanchard and Perotti (2002), BP henceforth. Given the recent reappraisal of fiscal policy, the question on the size of multipliers continues to be highly disputed (see e.g. Mountford and Uhlig, 2009; Ramey, 2011). BP estimate a three-variable VAR model of public spending, tax revenues and aggregate output. The SVAR model is identified through assumptions based on institutional knowledge: public spending does not respond to output in the quarter, while tax revenues do. Moreover, BP set the contemporaneous response of taxes to output on the basis of an outside estimate of the cyclical sensitivity of net taxes. Finally, they impose two alternative restrictions on the contemporaneous relationship between tax revenues and public spending, corresponding to two different models: in the first model a tax shock has an immediate effect on spending, but a spending shock does not have an immediate effect on tax revenues (except an indirect one through GDP), in the second model the other way around.

Despite the plausibility of the identification strategy, it is important to test such restrictions, at least for two reasons: (i) as Caldara and Kamps (2017) show, the use of plausible range of estimated elasticities may lead to dynamical responses and fiscal shocks that significantly differ in size and persistence; (ii) the authors are not able to distinguish the contemporaneous relationship between government spending and tax revenues and, consequently, whether a tax shock has an immediate effect on spending. Our empirical exercise is similar to the recent study by Karamysheva and Skrobotov (2022), who exploit the non-Gaussianity of the underlying reduced-form residuals to estimate and identify the same VAR as BP, using the GMM estimator derived in Lanne and Luoto (2021). Differently from this study, however, our exercise tests whether the identifying restrictions proposed by BP are rejected through our inference procedure. Moreover, we study to what extents results change when we use different ICA estimators, whose performance has been evaluated in the previous sections.

Table 2 shows the impact coefficients, derived in BP, under the unit normalization of the direct contemporaneous effects, i.e. the impact (mixing) matrix has been normalized so that it displays only ones on the main diagonal. Since, as mentioned, BP use two different models in function of the different contemporaneous impacts between spending

and taxes, we display both of them in the table. Note that, in terms of zero-entries, the only difference between the two models is the (1,2) entry. In the second model (right panel) the tax shock has a non-zero effect on spending but the spending shock has still a non-zero effect on tax because of the causal chain from  $G$  to  $Tax$  via  $GDP$ .

Table 2: Impact coefficients in Blanchard and Perotti (2002)

	G ordered first			Tax ordered first		
	$\varepsilon^g$	$\varepsilon^{tax}$	$\varepsilon^{gdp}$	$\varepsilon_t^g$	$\varepsilon_t^{tax}$	$\varepsilon^{gdp}$
$G$	1.00	0.00	0.00	1	-0.05	0.00
$Tax$	0.16	1.00	2.18	0.33	1.00	2.18
$GDP$	0.18	-0.15	1.00	0.15	-0.16	1

We estimate a three-variables VAR model using the same data as in BP (quarterly US data 1960-1997). The objective of the exercise is two-fold: (i) using our *Maxfinder* algorithm, to globally identify via the ICA model an independent shock in government spending and tax increase and estimate their dynamic responses on economic activity; (ii) to test the validity of restrictions proposed by BP, relying on the results on statistical inference obtained in Section 3.2.

For the analysis, we estimate a reduced-form VAR model analogous to equation (2), namely

$$y_t = C_t \Gamma + \sum_{l=1}^q B_l y_{t-l} + B_l^Q (Q1 + Q2 + Q3) y_{t-l} + u_t, \quad (12)$$

where  $y_t$  is the  $k \times 1$  vector of endogenous variables,  $C_t$  is  $k \times p$  vector of deterministic terms with coefficients  $\Gamma$  ( $p \times 1$ ),  $Q1$ ,  $Q2$ ,  $Q3$  are quarter dummies,  $B_l$  and  $B_l^Q$  are  $k \times k$  coefficient matrices, and  $u_t$  is  $k$ -dimensional vector of reduced-form residuals. In our case  $k = 3$  and  $q = 4$ , as in BP<sup>9</sup>. In line with the setting of the ICA model, we assume that the reduced-form residuals are a linear combination of statistically independent components, as specified in equation (3). As in BP, we have

$$y_t = (G_t, TAX_t, GDP_t)' \quad (13)$$

where  $y_t$  is a vector that contains the logarithm of real per capita values of government spending, taxes and GDP, observed in U.S. from 1960:Q1 to 1997:Q4. After estimation,

<sup>9</sup>The series of  $p$  deterministic terms include a constant, linear and quadratic time trend, quarter dummies plus a current and four-period lagged dummy for 1974:Q2 (when a large tax cut has been observed). When estimating the IRFs instead, closely following BP, we drop from the specification the non-linear terms  $Q \bullet y_{t-l}$ . The purpose of this exercise is to drop any serial correlation from the residuals, upon which the structural analysis is conducted.

the VAR is stable with no-serial autocorrelation. Moreover, as Figure 5 suggests, only for the spending's reduced-form residual,  $u_t^G$ , the Jarque-Bera test does not reject the null hypothesis of normality, so that ICA can be applied, since a single exception to non-Gaussianity is allowed in the ICA model.

Table 3 reports the estimates  $\tilde{b}_{ij}$  of the entries  $b_{ij}$  of the mixing matrix  $B_0$ , estimated with the four ICA methods and identified by applying our *MaxFinder* scheme. The estimates  $\tilde{b}_{ij}$  are the median of 500 estimates obtained by 500 bootstrap replications of the equation (12) model (through replications of  $u_t$ ), with confidence intervals derived as in Hall (1992):

$$CI(\alpha) = \left[ \tilde{b}_{ij} - q_{(1-\alpha/2)}^*, \tilde{b}_{ij} - q_{\alpha/2}^* \right] \quad (14)$$

where  $q_{(1-\alpha/2)}^*$  and  $q_{\alpha/2}^*$  are the  $(1 - \alpha/2)$  and  $\alpha/2$ -quantiles of empirical distribution of the root  $(\hat{b}_{ij}^* - \tilde{b}_{ij})$ , with  $(\hat{b}_{ij}^*)$  being the estimate of  $b_{ij}$  from the bootstrap replication). For sake of comparison, we normalize all the coefficients so that each structural innovation has a unit contemporaneous impact on the log of the variable that refers to.

For all estimators, the entries of the resulting mixing matrix do not satisfy the property for which the greatest value of each column lies in a distinct row. Therefore the *Maxfinder* is not able to deliver a permuted mixing matrix such that the off-diagonal elements are greater than main diagonal entries. However, the column permutation that penalizes low absolute values on the main diagonal (according to the LiNGAM criterion<sup>10</sup> described in Section 2.2) delivers an interesting result: we have identified two independent shocks that increase mostly taxes and government spending, and a third shock that increases substantially output and taxes. Moreover, this configuration is compatible with Table 2, which reports BP's findings and their shocks' labeling. Therefore, we can label the first ( $\varepsilon^1$ ) and the second ( $\varepsilon^2$ ) shock as a spending and tax shock, respectively. All the methods deliver a positive and significant impact coefficients of spending on GDP. Moreover, we get an estimate of a positive and statistically significant contemporaneous response of the third shock on tax revenues, but which is lower than BP's estimate of the GDP shock on tax revenues (except for *PML* that delivers a higher value). Finally, with the exception of *PML*, we get a non-significant impact of  $\varepsilon^3$  on  $G$ , which is consistent with one of the BP's zero restriction if  $\varepsilon^3$  is interpreted as the GDP shock. Overall, the ICA-estimated tax

---

<sup>10</sup>The application of the mere criterion of penalisation of small values (second step of the LiNGAM algorithm by Shimizu et al. (2006)) on the main diagonal does not involve the causal recursiveness assumption on which LiNGAM is based.



Table 3: Estimates of the contemporaneous impacts coefficients from ICA models. The entries of the mixing matrices are calculated using the bootstrap-median with the corresponding confidence intervals and are normalized so that each structural innovation has a unit contemporaneous impact on the log of the variable that refers to.

	<i>PML</i>			<i>fastICA</i>			<i>DCov</i>			<i>CvM</i>		
	$\varepsilon_1$	$\varepsilon_2$	$\varepsilon_3$	$\varepsilon_1$	$\varepsilon_2$	$\varepsilon_3$	$\varepsilon_1$	$\varepsilon_2$	$\varepsilon_3$	$\varepsilon_1$	$\varepsilon_2$	$\varepsilon_3$
$G_t$	1	-0.036***	-0.197***	1	0.006	-0.061	1	0.011	0.05	1	0.013	-0.745
$Tax_t$	0.955***	1	3.908***	0.642*	1	1.122***	0.442	1	1.146**	0.839*	1	1.393*
$GDP_t$	0.287***	-0.165***	1	0.31*	-0.018	1	0.227*	-0.019	1	0.574*	-0.01	1

shock ( $\varepsilon^2$ ) seems not to have a significant impact on both spending and GDP (except for the *PML* estimate). This result about spending is in tune with the zero-restriction of the model 1 in BP (zero impact from tax shock to  $G$ ). Given the results of our assessments in the previous section, we tend to rely more on the higher precision and the better empirical coverage of *PML* and *fastICA* estimators.

Having estimated the structural shocks and having identified the impact coefficients via ICA, we can now compute the impulse response functions and compare them with those implied by the BP model, shown in Figure 6. We first comment the impulse responses of a positive public spending shock, shown in Figure 7. The response of output is clearly positive and statistically significant, both at impact and within the first year. Figure 8 instead, shows the responses to an independent positive shock in tax revenues. All methods show that the effects of independent tax shocks are negative in the long run, which is a finding consistent with BP. However, all ICA-methods (except *PML*) show non-significant effect on the impact and in the short run, which is at odds with the finding by BP.

This illustrative exercise has shown that a purely data-driven identification procedure of VAR models is possible and, with careful modeling decisions, can lead to convincing conclusions based exclusively on statistical properties of the data. BP's identifying restrictions are plausible not only because of institutional knowledge and insights from economic theory: they are also present in the data and the ICA model supports them, at least as regards the zero restrictions of the BP's first model ( $G$  ordered first). On the contemporaneous relation between tax revenues and public spending, our results suggest that public spending is the first mover and its immediate impact on taxes is positive. In partial contrast to BP's results, our estimated impulse response functions suggest that a fiscal policy guided by public spending has a clearer effect on economic activity than a

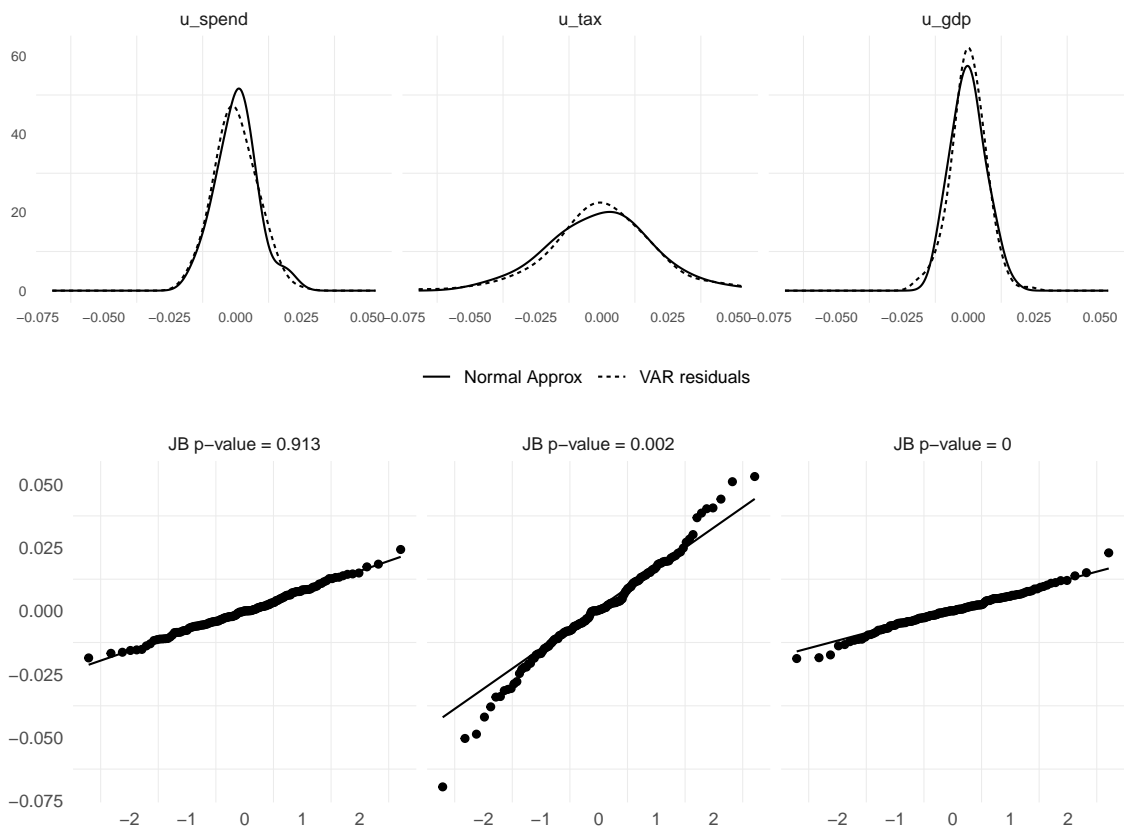


Figure 5: Distribution (upper panels) and q-q plots (bottom panels) of reduced-form residuals, where also p-values of the Jarque-Bera (JB) test are reported.

fiscal policy guided by tax revenues.

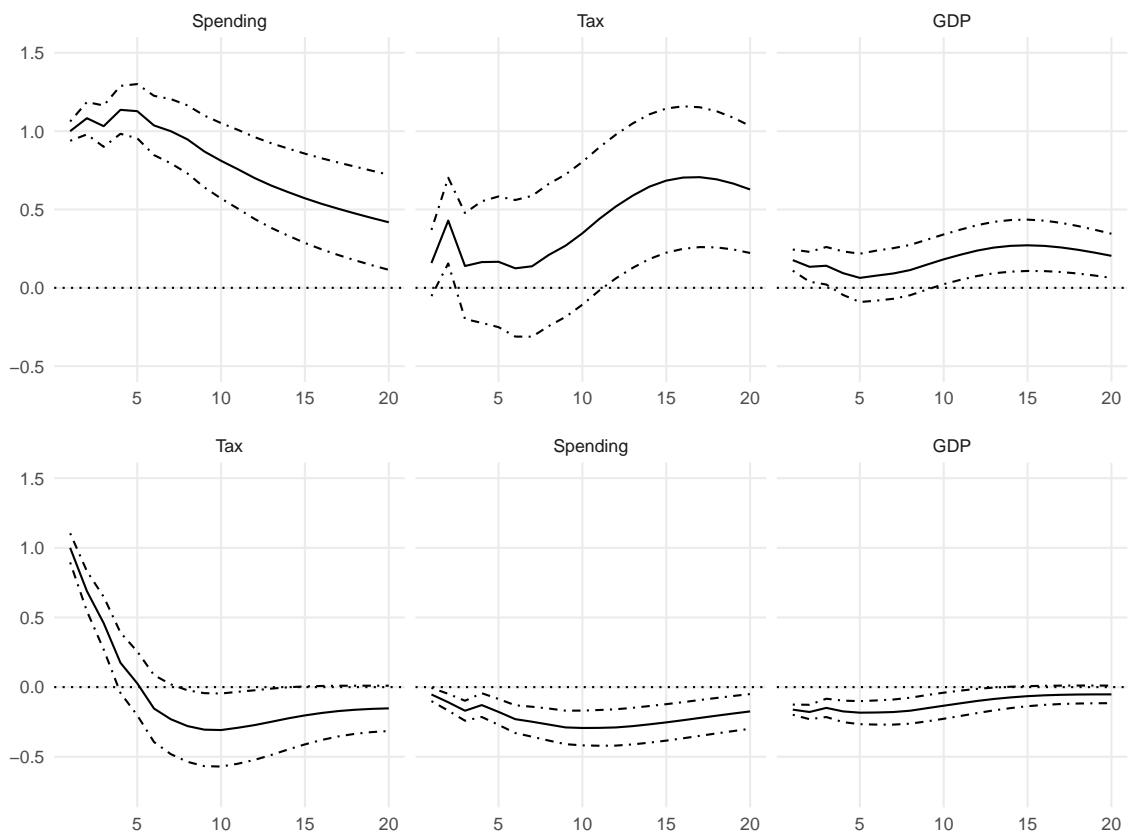


Figure 6: Impulse response functions estimated by Blanchard and Perotti (2002). The upper panel shows the responses of the three variables to a spending shock (on  $G$ ). The bottom panel shows the responses to a tax-revenue shock (on  $Tax$ ). Dashed lines denote an equal-tailed 68% confidence interval

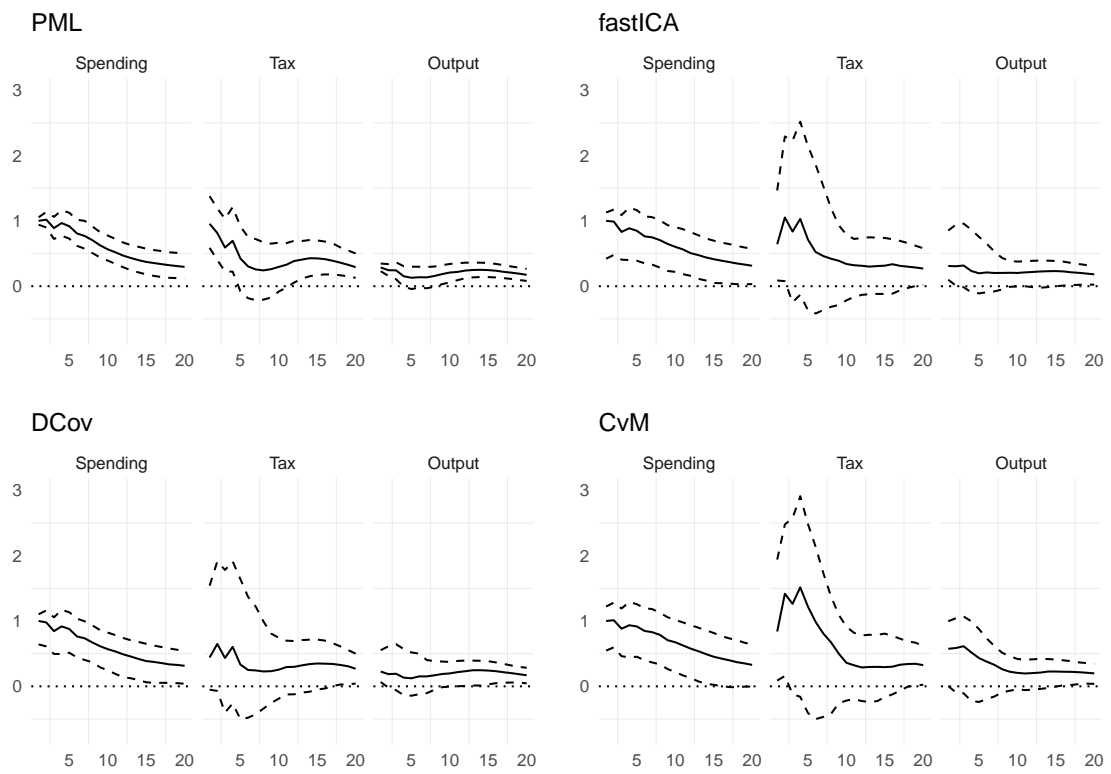


Figure 7: Impulse response functions of a positive public spending shock from different ICA estimation methods. The upper and the lower dashed lines represent respectively the 84% quantile and the 16% quantile of the bootstrap estimates.

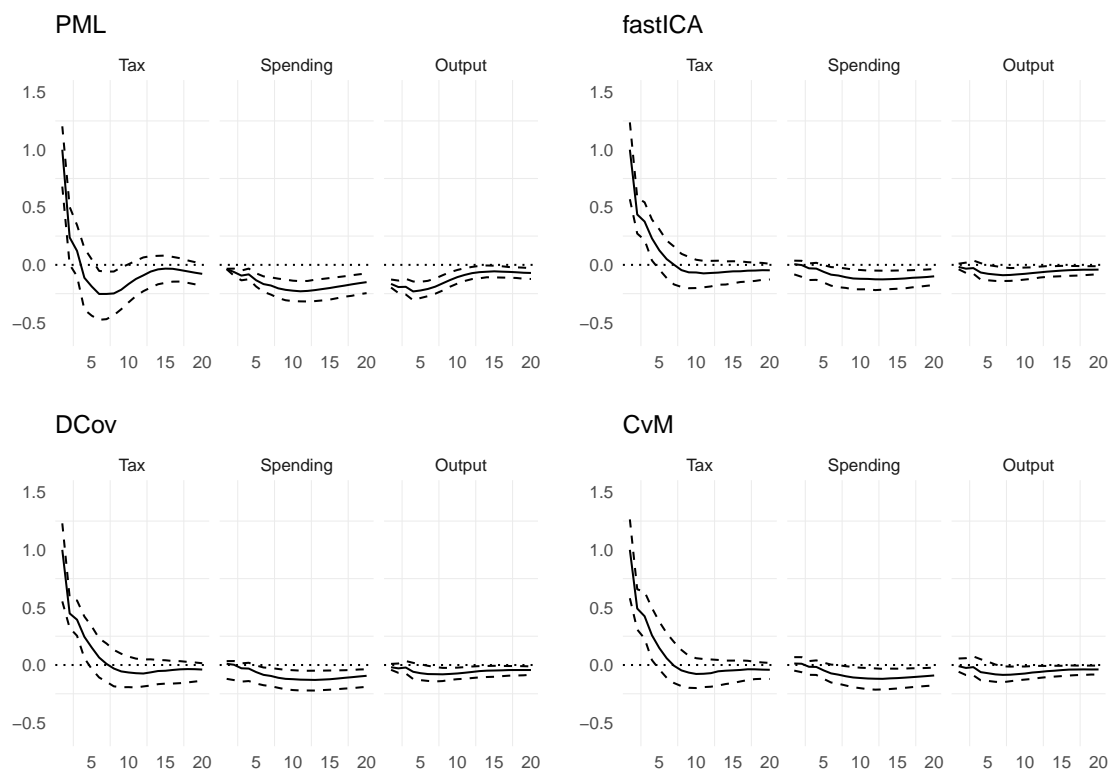


Figure 8: Impulse response functions of a positive tax shock from different ICA estimation methods. The upper and the lower dashed lines represent respectively the 84% quantile and the 16% quantile of the bootstrap estimates.

## 5 Conclusions

In this paper we assess, through Monte Carlo experiments, the performance of four ICA techniques (*fastICA*, *DCov*, *CvM*, *PML*), which have been recently used in SVAR analysis. We study the cases of structural disturbances that follow distributions which are both sub- and super-Gaussian but also approaching normality, that is the case in which the ICA model cannot recover the shocks by construction. The *fastICA* and the *PML* estimators result to be the ones showing a relatively lower variability and a more stable performance across both sub- and super-Gaussian settings, in comparison to the other estimators. In specific cases (e.g., number of variables  $k = 2$ ), the *DCov* estimator performs on average better than other estimators when the shocks' distributions are in the neighbourhood of Gaussianity. The variability of the *DCov* estimates, however, is relatively high and it is shared with the *CvM* estimator.

We also consider the distributions of the mixing matrix coefficients, which is the matrix that identifies a SVAR model and contains the simultaneous interactions of the variable of the system. Our Monte Carlo studies show that, as the dimensionality of the system increases, uncertainty in the estimates increases, as well as their negative bias. We also analyze the ICA methods' performance in statistical inference. Specifically, we have considered size distortions when testing the significance of the coefficients of the mixing matrix, comparing the performances of maximum likelihood versus bootstrap based inference. The *DCov* method, despite being relatively accurate on average, shows concerning variability. In the statistical inference exercise, on the other hand, the method based on the *PML* and *fastICA* estimators show lower size distortions and a better empirical coverage in almost all distributional scenarios. The *fastICA* estimator, more specifically, is the one which displays the lowest size distortion. Finally, an empirical application on fiscal policy highlights that a purely data-driven procedure such as ICA may help the researcher to test the significance of identifying restrictions or to suggest where to insert the latter. In particular, our exercise shows that the ICA model cannot reject the identification scheme implemented in Blanchard and Perotti (2002).

## References

- Acharya, D. and G. Panda (2008). A review of independent component analysis techniques and their applications. *IETE Technical Review* 25(6), 320–332.
- Back, A. D. and A. S. Weigend (1997). A first application of independent component analysis to extracting structure from stock returns. *International Journal of Neural Systems* 8(04), 473–484.
- Baqae, D. R. and E. Farhi (2019). The macroeconomic impact of microeconomic shocks: beyond Hulten’s theorem. *Econometrica* 87(4), 1155–1203.
- Bernanke, B. S. (1986). Alternative explanations of the money-income correlation. *Carnegie-Rochester Conference Series on Public Policy* 25, 49–99.
- Bernanke, B. S. and I. Mihov (1998). Measuring monetary policy. *The Quarterly Journal of Economics* 113(3), 869–902.
- Berner, A., S. Bruns, A. Moneta, and D. I. Stern (2022). Do energy efficiency improvements reduce energy use? Empirical evidence on the economy-wide rebound effect in Europe and the United States. *Energy Economics*, 105939.
- Blanchard, O. and R. Perotti (2002). An empirical characterization of the dynamic effects of changes in government spending and taxes on output. *Quarterly Journal of economics* 117(4), 1329–1368.
- Box, G. and D. Cox (1982). An analysis of transformations revisited, rebutted. *Journal of the American Statistical Association* 77(377), 209–210.
- Box, G. E. P. and G. C. Tiao (1962). A further look at robustness via Bayes’s theorem. *Biometrika* 49(3/4), 419–432.
- Brun, S. B., A. Moneta, and D. I. Stern (2021). Estimating the economy-wide rebound effect using empirically identified structural vector autoregressions. *Energy Economics* 97, 105158.
- Caldara, D. and C. Kamps (2017). The analytics of SVARs: a unified framework to measure fiscal multipliers. *The Review of Economic Studies* 84(3), 1015–1040.

- Capasso, M. and A. Moneta (2016). Macroeconomic responses to an independent monetary policy shock: A (more) agnostic identification procedure. Technical report, LEM Working Paper Series.
- Cardoso, J.-F. (1989). Source separation using higher order moments. In *International Conference on Acoustics, Speech, and Signal Processing*, pp. 2109–2112. IEEE.
- Ciarli, T., A. Coad, and A. Moneta (2019). Exporting and productivity as part of the growth process: Causal evidence from a data-driven structural VAR. Technical report, LEM Working Paper Series.
- Comon, P. (1994). Independent component analysis, a new concept? *Signal Processing* 36(3), 287–314.
- Eriksson, J. and V. Koivunen (2004). Identifiability, separability, and uniqueness of linear ICA models. *IEEE Signal Processing Letters* 11(7), 601–604.
- Fagiolo, G., M. Napoletano, and A. Roventini (2008). Are output growth-rate distributions fat-tailed? Some evidence from OECD countries. *Journal of Applied Econometrics* 23(5), 639–669.
- Fiorentini, G. and E. Sentana (2020). Discrete mixtures of normals pseudo maximum likelihood estimators of structural vector autoregressions. *CEPR Discussion Paper No. DP15411*.
- Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica* 79(3), 733–772.
- Genest, C., J.-F. Quessy, B. Rémillard, et al. (2007). Asymptotic local efficiency of Cramér–von Mises tests for multivariate independence. *The Annals of Statistics* 35(1), 166–191.
- Gertler, M. and P. Karadi (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics* 7(1), 44–76.
- Giacomini, R., D. N. Politis, and H. White (2013). A warp-speed method for conducting Monte Carlo experiments involving bootstrap estimators. *Econometric Theory* 29(3), 567–589.
- Goodman, I. R. and S. Kotz (1973). Multivariate  $\theta$ -generalized normal distributions. *Journal of Multivariate Analysis* 3(2), 204–219.



- Gouriéroux, C., A. Monfort, and J.-P. Renne (2017). Statistical inference for independent component analysis: Application to structural VAR models. *Journal of Econometrics* 196(1), 111–126.
- Gouriéroux, C., A. Monfort, and J. P. Renne (2020). Identification and estimation in non-fundamental structural VARMA models. *The Review of Economic Studies* 87(4), 1915–1953.
- Guay, A. (2021). Identification of structural vector autoregressions through higher unconditional moments. *Journal of Econometrics* 225(1), 27–46.
- Guerini, M., A. Moneta, M. Napoletano, and A. Roventini (2020). The janus-faced nature of debt: Results from a data-driven cointegrated SVAR approach. *Macroeconomic Dynamics* 24(1), 24–54.
- Hall, P. (1992). Effect of bias estimation on coverage accuracy of bootstrap confidence intervals for a probability density. *The Annals of Statistics* 20(2), 675–694.
- Hastie, T. and R. Tibshirani (2010). *ProDenICA: Product Density Estimation for ICA using tilted Gaussian density estimates*. R Foundation for Statistical Computing. R package version 1.0.
- Herwartz, H. (2018). Hodges–Lehmann detection of structural shocks – An analysis of macroeconomic dynamics in the Euro area. *Oxford Bulletin of Economics and Statistics* 80(4), 736–754.
- Herwartz, H. (2019). Long-run neutrality of demand shocks: Revisiting Blanchard and Quah (1989) with independent structural shocks. *Journal of Applied Econometrics* 34(5), 811–819.
- Herwartz, H., A. Lange, and S. Maxand (2021). Data-driven identification in SVARs—When and how can statistical characteristics be used to unravel causal relationships? *Economic Inquiry*.
- Herwartz, H. and M. Plödt (2016). The macroeconomic effects of oil price shocks: Evidence from a statistical identification approach. *Journal of International Money and Finance* 61, 30–44.
- Horn, R. A. and C. R. Johnson (2012). *Matrix analysis*. Cambridge University Press.

- Hyvärinen, A. (1999). Fast and robust fixed-point algorithms for independent component analysis. *IEEE transactions on Neural Networks* 10(3), 626–634.
- Hyvärinen, A. (2013). Independent component analysis: recent advances. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 371(1984), 20110534.
- Hyvärinen, A. and E. Oja (1997). A fast fixed-point algorithm for independent component analysis. *Neural computation* 9(7), 1483–1492.
- Hyvärinen, A. and E. Oja (2000). Independent component analysis: algorithms and applications. *Neural networks* 13(4-5), 411–430.
- Hyvärinen, A., K. Zhang, S. Shimizu, and P. O. Hoyer (2010). Estimation of a structural vector autoregression model using non-Gaussianity. *Journal of Machine Learning Research* 11(56), 1709–1731.
- Ilmonen, P., K. Nordhausen, H. Oja, and E. Ollila (2010). A new performance index for ICA: properties, computation and asymptotic analysis. In *International Conference on Latent Variable Analysis and Signal Separation*, pp. 229–236. Springer.
- Kalke, S. (2015). *pgnorm: The p-Generalized Normal Distribution*. R Foundation for Statistical Computing. R package version 2.0.
- Kalke, S. and W.-D. Richter (2013). Simulation of the p-generalized Gaussian distribution. *Journal of Statistical Computation and Simulation* 83(4), 641–667.
- Karamysheva, M. and A. Skrobotov (2022). Do we reject restrictions identifying fiscal shocks? Identification based on non-Gaussian innovations. *Journal of Economic Dynamics and Control* 138, 104358.
- Kilian, L. and H. Lütkepohl (2017). *Structural Vector Autoregressive analysis*. Cambridge University Press.
- Klein, L. R. and A. S. Goldberger (1955). *An Econometric Model of the United States, 1929-1952*. North Holland.
- Lacerda, G., P. Spirtes, J. Ramsey, and P. O. Hoyer (2008). Discovering cyclic causal models by independent components analysis. In *Proceedings of the Twenty-Fourth Conference on*

- Uncertainty in Artificial Intelligence*, UAI 08, Arlington, Virginia, USA, pp. 366–374. AUAI Press.
- Lanne, M. and J. Luoto (2021). GMM estimation of non-Gaussian structural vector autoregression. *Journal of Business & Economic Statistics* 39(1), 69–81.
- Lanne, M. and H. Lütkepohl (2010). Structural vector autoregressions with nonnormal residuals. *Journal of Business & Economic Statistics* 28(1), 159–168.
- Lanne, M., M. Meitz, and P. Saikkonen (2017). Identification and estimation of non-Gaussian structural vector autoregressions. *Journal of Econometrics* 196(2), 288–304.
- Lütkepohl, H. and F. Xu (2012, Jun). The role of the log transformation in forecasting economic variables. *Empirical Economics* 42(3), 619–638.
- Matteson, D. S. and R. S. Tsay (2017). Independent component analysis via distance covariance. *Journal of the American Statistical Association* 112(518), 623–637.
- Maxand, S. (2020). Identification of independent structural shocks in the presence of multiple gaussian components. *Econometrics and Statistics* 16, 55–68.
- Moneta, A., D. Entner, P. O. Hoyer, and A. Coad (2013). Causal inference by independent component analysis: Theory and applications. *Oxford Bulletin of Economics and Statistics* 75(5), 705–730.
- Mountford, A. and H. Uhlig (2009). What are the effects of fiscal policy shocks? *Journal of Applied Econometrics* 24(6), 960–992.
- Nelson, H. L. and C. Granger (1979). Experience with using the box-cox transformation when forecasting economic time series. *Journal of Econometrics* 10(1), 57–69.
- Nordhausen, K., E. Ollila, and H. Oja (2011). On the performance indices of ICA and blind source separation. In *Signal Processing Advances in Wireless Communications (SPAWC), 2011 IEEE 12th International Workshop on*, pp. 486–490. IEEE.
- Ramey, V. A. (2011). Identifying government spending shocks: It’s all in the timing. *Quarterly Journal of Economics* 126(1), 1–50.
- Rigobon, R. (2003). Identification through heteroskedasticity. *Review of Economics and Statistics* 85(4), 777–792.

- Romer, C. D. and D. H. Romer (2010). The macroeconomic effects of tax changes: estimates based on a new measure of fiscal shocks. *American Economic Review* 100(3), 763–801.
- Shannon, C. E. (1949). *The Mathematical Theory of Communication*. By CE Shannon and Warren Weaver. Urbana.
- Shimizu, S., P. O. Hoyer, A. Hyvärinen, and A. Kerminen (2006). A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research* 7(Oct), 2003–2030.
- Sims, C. A. (1980). Macroeconomics and reality. *Econometrica* 48(1), 1–48.
- Sims, C. A. (1986). Are forecasting models usable for policy analysis? *Quarterly Review* 10, 2–16.
- Stock, J. H. and M. W. Watson (2001). Vector autoregressions. *Journal of Economic Perspectives* 15(4), 101–115.
- Stock, J. H. and M. W. Watson (2016). Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics. In *Handbook of macroeconomics*, Volume 2, pp. 415–525. Elsevier.
- Stock, J. H. and M. W. Watson (2017). Twenty years of time series econometrics in ten pictures. *Journal of Economic Perspectives* 31(2), 59–86.
- Subbotin, M. T. (1923). On the law of frequency of error. *Math. USSR-Sb* 31(2), 296–301.
- Székely, G. J., M. L. Rizzo, N. K. Bakirov, et al. (2007). Measuring and testing dependence by correlation of distances. *Annals of Statistics* 35(6), 2769–2794.
- Tiao, G. and D. R. Lund (1970). The use of OLUMV estimators in inference robustness studies of the location parameter of a class of symmetric distributions. *Journal of the American Statistical Association* 65(329), 370–386.
- 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.
- Yu, S., A. Zhang, and H. Li (2012). A review of estimating the shape parameter of generalized Gaussian distribution. *J. Comput. Inf. Syst* 8(21), 9055–9064.
- Zema, S. M. (2021). Directed acyclic graph based information shares for price discovery. Technical report, SSRN.

## Appendix A. Other Figures

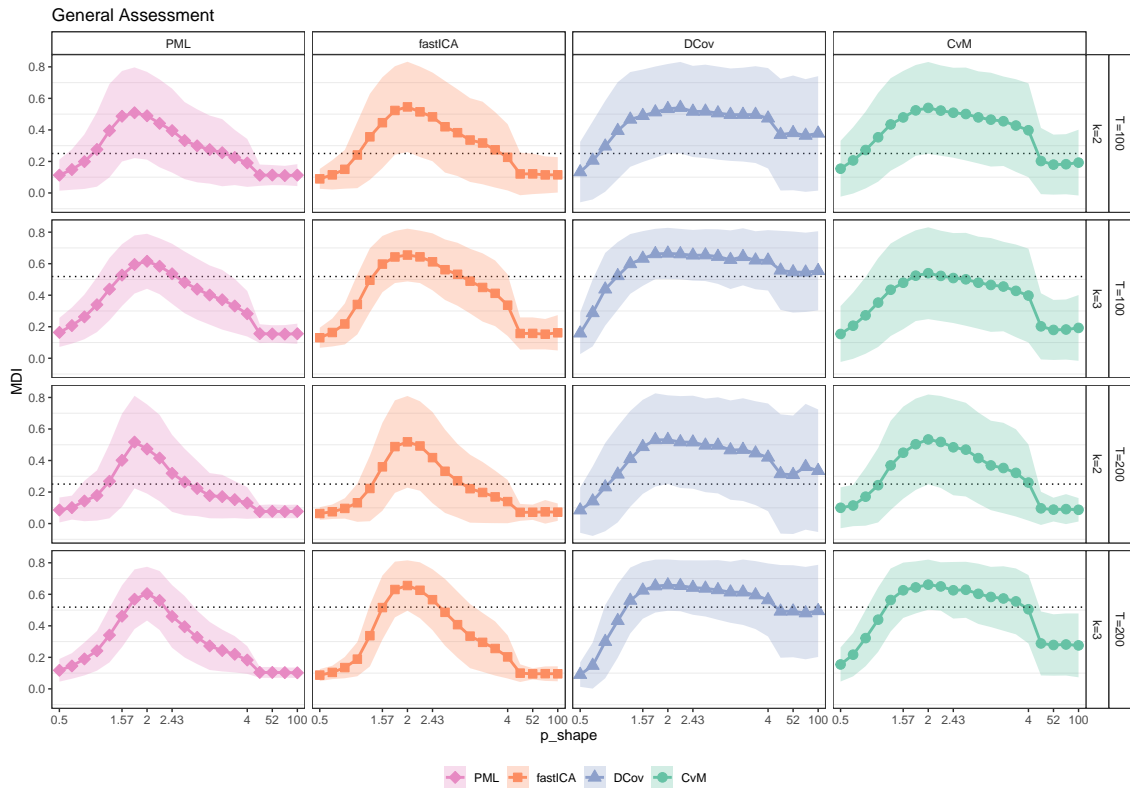


Figure 9: Each plot shows the MDI (thick lines with marks: averages across 1000 dgps), as specified in equation (7), for each of the four ICA estimators, across 20 values of  $p$  (*general assessment*). The  $p$ -generalized normal distribution is super-Gaussian with  $p < 2$ , Gaussian with  $p = 2$ , sub-Gaussian with  $p > 2$ . Plots in the upper part of the figure correspond to the case  $k = 2$ , in the lower part to  $k = 3$ . Sample size  $T = 400$ . Shadow areas show one-standard deviation above and below the mean. Dashed lines constitute the negative benchmark: 90% of random estimates yield MDIs above this line.

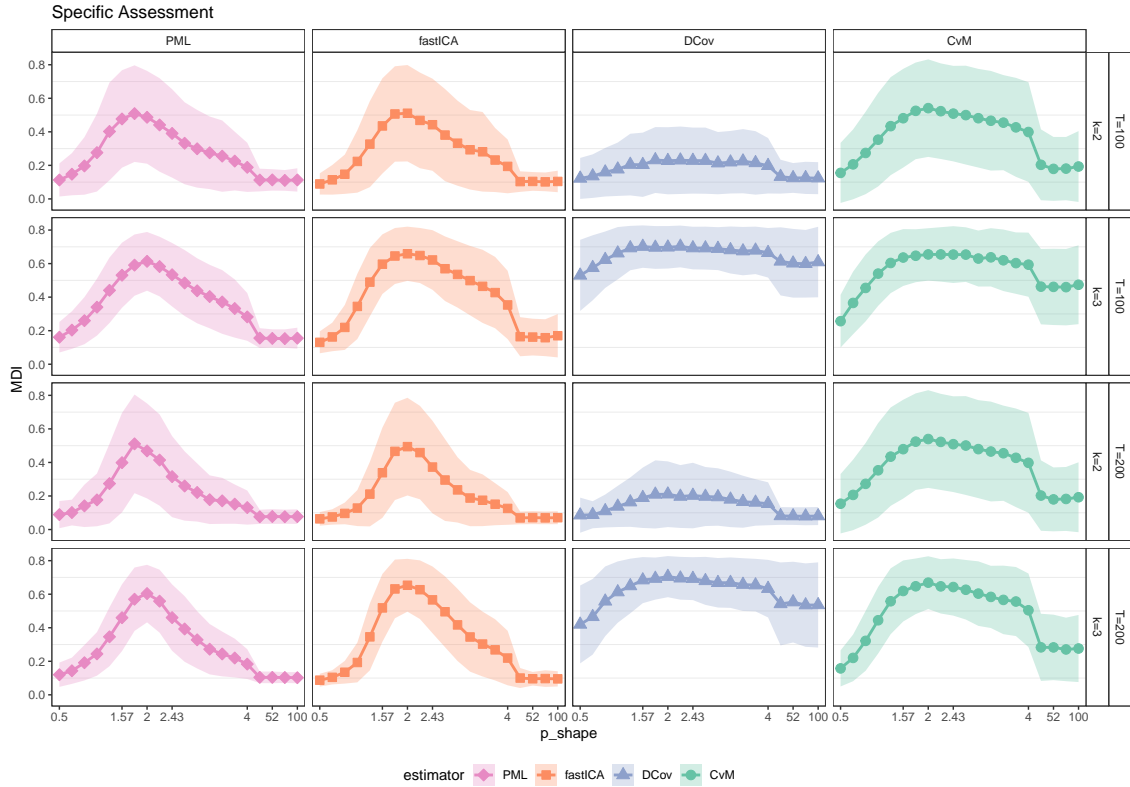


Figure 10: Each plot shows the MDI (thick lines with marks: averages across 1000 Monte Carlo replications), as specified in equation (9) (*specific assessment*), for each of the four ICA estimators, across 20 values of  $p$ . The mixing matrix in the  $dgp$  is fixed as in equation (8). The  $p$ -generalized normal distribution is super-Gaussian with  $p < 2$ , Gaussian with  $p = 2$ , sub-Gaussian with  $p > 2$ . Plots in the upper part of the figure correspond to the case  $k = 2$ , in the lower part to  $k = 3$ . Sample size  $T = 400$ . Shadow areas show one-standard deviation above and below the mean.

## Appendix B. R Codes and Data

The R codes and data to replicate the results reported in this article are available at the following GitHub repository:

<https://github.com/gianlucaallante/ica-svars-comparative-analysis>