

SVAR Identification From Higher Moments: Has the Simultaneous Causality Problem Been Solved?

By JOSÉ LUIS MONTIEL OLEA, MIKKEL PLAGBORG-MØLLER, AND ERIC QIAN*

Two recent strands of the literature on Structural Vector Autoregressions (SVARs) use higher moments for identification. One of them exploits independence and non-Gaussianity of the shocks; the other, stochastic volatility (heteroskedasticity). These approaches achieve point identification without imposing exclusion or sign restrictions. We review this work critically, and contrast its goals with the separate research program that has pushed for macroeconomics to rely more heavily on credible economic restrictions and institutional knowledge, as is the standard in microeconomic policy evaluation. Identification based on higher moments imposes substantively stronger assumptions on the shock process than standard second-order SVAR identification methods do. We recommend that these assumptions be tested in applied work. Even when the assumptions are not rejected, inference based on higher moments necessarily demands more from a finite sample than standard approaches do. Thus, in our view, weak identification issues should be given high priority by applied users.

I. Identification From Second Moments

We first review the well-known issues in identification of SVARs from second moments of the data. Suppose the data $\mathbf{y}_t = (y_{1,t}, y_{2,t})'$ follows the SVAR process

$$(1) \quad \mathbf{y}_t = \mathbf{c} + \sum_{\ell=1}^p \mathbf{A}_\ell \mathbf{y}_{t-\ell} + \mathbf{H} \boldsymbol{\varepsilon}_t,$$

* Montiel Olea: Department of Economics, Columbia University, 1022 International Affairs Building New York, NY 10027, montiel.olea@gmail.com. Plagborg-Møller: Department of Economics, Princeton University, Julius Romo Rabinowitz Building, Princeton, NJ 08544, mikkelpm@princeton.edu. Qian: Department of Economics, Princeton University, Julius Romo Rabinowitz Building, Princeton, NJ 08544, eric-qian@princeton.edu. We are grateful for comments by Thorsten Drautzburg, Jim Hamilton, Daniel Lewis, Helmut Lutkepohl, Geert Meesters, Sophocles Mavroidis, Enrique Sentana, and Christian Wolf. Plagborg-Møller acknowledges support from the NSF under Grant #1851665.

where the vector of unobserved shocks $\boldsymbol{\varepsilon}_t = (\varepsilon_{1,t}, \varepsilon_{2,t})'$ is assumed to be orthogonal white noise, that is, serially and mutually uncorrelated: $\text{Cov}(\boldsymbol{\varepsilon}_t, \boldsymbol{\varepsilon}_{t-\ell}) = 0$ for $\ell \geq 1$ and $\text{Cov}(\varepsilon_{1,t}, \varepsilon_{2,t}) = 0$. The diagonal entries of $\mathbf{H}$ are normalized to 1.

Due to the familiar simultaneous causality problem, the basic SVAR model is not identified. The autoregressive matrices $\mathbf{A}_\ell$ in (1) can be identified through a projection of $\mathbf{y}_t$ on its lags (assuming stationarity); let $\boldsymbol{\eta}_t = (\eta_{1,t}, \eta_{2,t})' = \mathbf{H} \boldsymbol{\varepsilon}_t$ be the reduced-form residuals from this projection. If we knew $\mathbf{H}$, we could identify the dynamic effects of the shocks $\boldsymbol{\varepsilon}_t = \mathbf{H}^{-1} \boldsymbol{\eta}_t$. Unfortunately, it is not possible to identify $\mathbf{H}$ without further assumptions. Since the white noise assumption on the shocks only has implications for first and second moments of the data, the only exploitable information about $\mathbf{H}$ comes from the equation

$$(2) \quad \text{Var}(\boldsymbol{\eta}_t) = \mathbf{H} \begin{pmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{pmatrix} \mathbf{H}',$$

where $\sigma_j^2 = \text{Var}(\varepsilon_{j,t})$. The non-linear Equation (2) has multiple solutions. Intuitively, the left-hand side of (2) only has three non-redundant (consistently estimable) reduced-form parameters due to symmetry, while the right-hand side has four unknowns (the off-diagonal elements of $\mathbf{H}$ and variances of the shocks).

To overcome the simultaneous causality issue, a vast literature on SVAR identification has proposed to exploit various additional economic restrictions, such as exclusion or sign restrictions. These restrictions are usually motivated using economic theory and/or the institutional background for the application at hand. For example, in recent years, it has become popular to achieve identification by constructing instrumental variables (also known as proxies) that are assumed to correlate only with the shock of interest but not other shocks. The practice of explicitly defending these instrument exclu-

sion restrictions brings identification in macroeconomics closer to the *modus operandi* in applied microeconomics (Nakamura and Steinson, 2018*b*; Stock and Watson, 2018).

SVAR estimators based on second moments are usually highly robust to statistical properties of the data. For example, the usual Ordinary Least Squares (OLS) estimator can be viewed as the quasi-maximum likelihood estimator (MLE) of the model (1) under the working assumptions that the shocks ε_t are i.i.d., Gaussian, and homoskedastic. But importantly, none of these working assumptions are necessary for the consistency of the OLS estimator (e.g., Gonçalves and Kilian, 2004).¹

II. Identification From Mutual Independence and Non-Gaussianity

Beginning with Gouriéroux, Monfort and Renne (2017) and Lanne, Meitz and Saikkonen (2017), several recent papers have achieved SVAR identification by importing techniques from the statistics literature on Independent Components Analysis (ICA, see the review by Hyvärinen, Karhunen and Oja, 2001). The key assumptions in this literature are that the shocks $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$ are mutually independent and non-Gaussian. Given these assumptions, it is natural to ask:

Question: If we find two linear combinations $C_{11}\eta_{1,t} + C_{12}\eta_{2,t}$ and $C_{21}\eta_{1,t} + C_{22}\eta_{2,t}$ of the reduced-form residuals that are mutually independent, can we conclude that these linear combinations must equal the true shocks $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$ (up to scale and ordering)?

Answer, part 1: If the shocks were Gaussian, it is well-known that the answer is *no*. There exist multiple linear combinations that are uncorrelated and, due to Gaussianity, therefore also independent.

Answer, part 2: If at least one of the shocks $\varepsilon_{1,t}$ or $\varepsilon_{2,t}$ is *not* Gaussian, then the answer is *yes*. This follows from the Darmois-Skitovich Theorem (Comon, 1994): Two non-trivial linear combinations of independent random variables cannot themselves be independent, unless the variables are all Gaussian.

¹Another attractive feature of the Gaussian quasi-MLE is that the asymptotic limit of the resulting impulse response estimates can be interpreted as arising from Local Projections with a transparent interpretation (Plagborg-Møller and Wolf, 2021).

Intuitively, though second moments of the data do not suffice for identification (cf. Section I), higher moments provide further restrictions to pin down the structural parameters. Third or higher moments are redundant if the shocks are exactly Gaussian, but since there is no reason to believe that this is the case in practice, the identification result appears very attractive.

Following the ICA literature, several SVAR estimators based on non-Gaussian identification are available, such as estimators that exploit restrictions on third and fourth moments implied by independence (Lanne and Luoto, 2021) or quasi-MLE (Gouriéroux, Monfort and Renne, 2017; Lanne, Meitz and Saikkonen, 2017; Fiorentini and Sentana, 2020; Sims, 2021). Many of these estimators are semiparametrically consistent, i.e., without requiring the functional form of the shock distribution to be known.

As is clear, the SVAR-ICA approach overcomes the simultaneous causality problem, not by exploiting economic theory or institutional details, but by imposing the strong statistical assumption that the shocks are mutually independent.² This substantially strengthens the orthogonal white noise assumption in standard SVAR analysis. For example, it rules out the white noise shock process

$$(3) \quad \varepsilon_{j,t} = \tau_t \zeta_{j,t}, \quad j = 1, 2,$$

where τ_t is a shared scalar volatility process that is independent of the i.i.d., independent, and mean-zero processes $\zeta_{1,t}$ and $\zeta_{2,t}$. This process is consistent with historical episodes such as the Great Moderation and the Covid-19 recession.

One could argue that the independence assumption is innocuous because it can be justified from the definition of structural shocks in the impulse-propagation paradigm for macroeconomic dynamics. While this is a reasonable argument, the SVAR-ICA literature further assumes that these independent shocks enter additively in the SVAR model (1). To spell this out in an example, consider the SVAR model with stochastic volatility in Equations (1) and (3). One could argue that the basic shocks in this model are $\zeta_{1,t}$ and $\zeta_{2,t}$ (as well as shocks

²A modicum of economic reasoning is needed to label the estimated shocks, for example by inspecting impulse responses.

to τ_t), which are indeed mutually independent. Traditional linear SVAR approaches based on second moments will not estimate impulse responses with respect to $\zeta_{1,t}$ and $\zeta_{2,t}$, but if their identifying restrictions are satisfied, they will consistently estimate impulse responses with respect to the orthogonal white noise disturbances $\varepsilon_{1,t} = \tau_t \zeta_{1,t}$ and $\varepsilon_{2,t} = \tau_t \zeta_{2,t}$, which are arguably still interesting structural objects. In contrast, since $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$ are not independent in this example, ICA approaches may fail to estimate any object with a structural interpretation.³

If one insists on estimating impulse responses with respect to independent shocks, it seems necessary to consider classes of non-linear models that—at a minimum—are rich enough to allow for shared (and persistent) volatility factors in the residuals, such as (3). Nonlinear variants of ICA exist (Hyvärinen, Karhunen and Oja, 2001, Ch. 17), but these have not yet been adapted to macroeconomic applications.

We stress that the choice between identification based on second moments vs. higher-order moments is *not* a choice between Gaussian shocks vs. non-Gaussian shocks. Second-moment SVAR estimators require neither Gaussianity, nor mutual shock independence, as discussed in Section I.

III. Testing Mutual Independence

As pointed out by Matteson and Tsay (2017), Amengual, Fiorentini and Sentana (2021), and Davis and Ng (2021), the assumption of mutual shock independence is testable in the data. Given any ICA estimator $\hat{\mathbf{H}}$ that is consistent under independence, one simply needs to check whether the estimated shocks $(\hat{\varepsilon}_{1,t}, \hat{\varepsilon}_{2,t})' = \hat{\mathbf{H}}^{-1} \hat{\boldsymbol{\eta}}_t$ are in fact independent of each other. In the Online Appendix we implement a particular version of this idea, using a bootstrap to account for estimation error. We reject independence in some specifications of a canonical monetary SVAR model applied to U.S. data, though not in all specifications. This suggests that shock independence should not be viewed as an unobjec-

tionable assumption, but rather as a hypothesis to be tested in every application. It is an interesting area of research to further develop and analyze such tests.⁴

IV. Identification From Heteroskedasticity

A different approach to higher-moment identification exploits heteroskedasticity (i.e., time-varying conditional volatility) in the data, following Sentana and Fiorentini (2001) and Rigobon (2003). See Lewis (2021a) and Sims (2021) for recent results and references. Instead of assuming serial and mutual shock independence, these papers make the weaker assumption that the shocks are non-linearly unpredictable and conditionally (not just unconditionally) orthogonal:

$$E(\varepsilon_t | \mathcal{I}_{t-1}) = 0, \text{Cov}(\varepsilon_{1,t}, \varepsilon_{2,t} | \mathcal{I}_{t-1}) = 0,$$

for all t , where $\mathcal{I}_{t-1}$ is an information set available to the econometrician at time $t - 1$. Denote the conditional variance of shock j by $\sigma_{j,t-1}^2 = \text{Var}(\varepsilon_{j,t} | \mathcal{I}_{t-1})$. Then, similar to Equation (2),

$$\text{Var}(\boldsymbol{\eta}_t | \mathcal{I}_{t-1}) = \mathbf{H} \begin{pmatrix} \sigma_{1,t-1}^2 & 0 \\ 0 & \sigma_{2,t-1}^2 \end{pmatrix} \mathbf{H}'.$$

Denote the left-hand side by Σ_t ; note that this conditional reduced-form variance-covariance matrix is in principle estimable in the data. Then it follows that

$$(4) \quad \Sigma_t \Sigma_{t-1}^{-1} = \mathbf{H} \begin{pmatrix} \frac{\sigma_{1,t}^2}{\sigma_{1,t-1}^2} & 0 \\ 0 & \frac{\sigma_{2,t}^2}{\sigma_{2,t-1}^2} \end{pmatrix} \mathbf{H}^{-1}.$$

In words, the columns of the matrix $\mathbf{H}$ are the eigenvectors of the matrix $\Sigma_t \Sigma_{t-1}^{-1}$. These eigenvectors are uniquely determined (given the normalization $H_{11} = H_{22} = 1$), provided that the eigenvalues $\frac{\sigma_{1,t}^2}{\sigma_{1,t-1}^2}$ and $\frac{\sigma_{2,t}^2}{\sigma_{2,t-1}^2}$ are distinct, i.e., if the volatilities of the two shocks jump by different amounts from time $t - 1$ to time t . This rules out the model (3) with a single shared volatility factor, for example.

³For example, if the $\zeta_{j,t}$'s are standard normal, then ε_t has a non-normal but spherical distribution. Hence, in population, the Gouriéroux, Monfort and Renne (2017) quasi-likelihood is a constant function of its argument, the orthogonal rotation matrix. In sample, this implies that the quasi-MLE will be an essentially arbitrary rotation of the true impulse responses.

⁴Separately, one could test the “relevance condition” that the shocks are non-normal using conventional normality tests (e.g., Lanne, Meitz and Saikkonen, 2017, Sec. 5.2).

In summary, if we impose the stronger assumption that the shocks are conditionally orthogonal, we are able to identify the SVAR model by essentially exploiting conditional second moments in multiple volatility “regimes”. This is still a “higher-moment” procedure in the sense of using information beyond unconditional second moments.

We stress, once again, that the choice is *not* between assuming heteroskedasticity or homoskedasticity. Traditional (unconditional) second-moment procedures often use a homoskedastic Gaussian quasi-MLE, but homoskedasticity is not required for either consistency or identification.

The above identification argument immediately suggests a way to test the identifying assumption of conditional shock orthogonality: Check whether the eigenvectors of $\Sigma_t \Sigma_{t-1}^{-1}$ are indeed constant at each point in time t (only the eigenvalues may vary), as predicted by Equation (4).⁵ This general idea applies to many different volatility models that have been considered in the literature, such as discrete Markov-switching volatility jumps and continuous volatility changes (Lewis, 2021a).

Popular SVAR procedures that exploit heteroskedasticity rely on assumptions that are completely at odds with popular SVAR-ICA procedures. Whereas the former estimate time-varying volatilities, the latter usually assume that shocks are i.i.d. over time (though this is not required for identification). There does not appear to be any reason why the two different kinds of estimation procedures could not be made mutually consistent. We hope that future research will pursue this, while addressing the shortcomings we have mentioned previously.

V. Sensitivity of Higher-Moment Identification

Any identification approach that relies on third or higher moments is at risk of suffering from weak identification. Recall that these approaches must necessarily fail when the data is (unconditionally) Gaussian, since then higher moments provide no additional information. It is of no comfort that macroeconomic data is unlikely to be exactly Gaussian; the relevant ques-

tion is whether the data generating process is sufficiently non-Gaussian relative to the estimation error in the higher moments. As is well known, the estimation error in higher moments is typically large for moderate sample sizes.

The potential for weak identification suggests that higher-moment procedures could be more sensitive to minor perturbations of the data than conventional second-moment procedures are. Indeed, Lanne and Luoto (2021, Sec. 3.4) demonstrate in a model with t-distributed shocks that the performance of SVAR-ICA estimators and tests deteriorates as the degrees of freedom increase.

We therefore think it is important for the literature on higher-moment identification to discuss more seriously potential weak identification issues. At a minimum, applied researchers could easily report empirically calibrated simulations to verify that the higher-moment procedures behave as expected, given the actual sample size and potential non-Gaussianity of the residuals. Ideally, in the long-run, any reported inference results would be complemented by weak-identification-robust procedures, such as those that have been recently developed both for identification by ICA (Drautzburg and Wright, 2021; Lee and Mesters, 2021) and by heteroskedasticity (Nakamura and Steinsson, 2018a, Appendix C; Lewis, 2021b).

VI. Conclusion

This paper has argued the following. First, identification from higher moments does not simply exploit more information in the data than traditional SVAR methods; it requires stronger assumptions on the shock process than second-moment methods do. Second, applied work should routinely test the additional shock assumptions. It should also give high priority to issues of sensitivity and weak identification, since these are likely more acute when exploiting higher moments in macroeconomic data. Third, second-moment identification remains relevant due to its robustness to statistical properties of the data. Moreover, we believe it is a virtue rather than a limitation that researchers are forced to defend their identifying restrictions based on substantive economic or institutional knowledge, rather than appealing to statistical assumptions about the shock process.

⁵This is related to Rigobon’s (2003) over-identification test. Separately, one could test the “relevance condition” that the eigenvalues in (4) are distinct; see references in Lewis (2021a).

REFERENCES

- Amengual, Dante, Gabriele Fiorentini, and Enrique Sentana.** 2021. "Moment tests of independent components." *SERIEs*.
- Comon, Pierre.** 1994. "Independent component analysis, a new concept?" *Signal Processing*, 36: 287–314.
- Davis, Richard, and Serena Ng.** 2021. "Time Series Estimation of the Dynamic Effects of Disaster-Type Shocks." Manuscript, Columbia University.
- Drautzborg, Thorsten, and Jonathan H. Wright.** 2021. "Refining Set-Identification in VARs through Independence." NBER Working Paper 29316.
- Fiorentini, Gabriele, and Enrique Sentana.** 2020. "Discrete mixtures of normals pseudo maximum likelihood estimators of structural vector autoregressions." CEPR Discussion Paper DP15411.
- Gonçalves, Sílvia, and Lutz Kilian.** 2004. "Bootstrapping autoregressions with conditional heteroskedasticity of unknown form." *Journal of Econometrics*, 123(1): 89–120.
- Gouriéroux, Christian, Alain Monfort, and Jean-Paul Renne.** 2017. "Statistical inference for independent component analysis: Application to structural VAR models." *Journal of Econometrics*, 196(1): 111–126.
- Hyvärinen, Aapo, Juha Karhunen, and Erkki Oja.** 2001. *Independent Component Analysis*. John Wiley & Sons, Inc.
- Lanne, Markku, and Jani Luoto.** 2021. "GMM Estimation of Non-Gaussian Structural Vector Autoregression." *Journal of Business & Economic Statistics*, 39(1): 69–81.
- Lanne, Markku, Mika Meitz, and Pentti Saikkonen.** 2017. "Identification and estimation of non-Gaussian structural vector autoregressions." *Journal of Econometrics*, 196(2): 288–304.
- Lee, Adam, and Geert Mesters.** 2021. "Robust Non-Gaussian Inference for Linear Simultaneous Equations Models." Manuscript, Universitat Pompeu Fabra.
- Lewis, Daniel J.** 2021a. "Identifying Shocks via Time-Varying Volatility." *Review of Economic Studies*. Forthcoming.
- Lewis, Daniel J.** 2021b. "Robust Inference in Models Identified via Heteroskedasticity." *Review of Economics and Statistics*. Forthcoming.
- Matteson, David S., and Ruey S. Tsay.** 2017. "Independent Component Analysis via Distance Covariance." *Journal of the American Statistical Association*, 112(518): 623–637.
- Nakamura, Emi, and Jón Steinsson.** 2018a. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *The Quarterly Journal of Economics*, 133(3): 1283–1330.
- Nakamura, Emi, and Jón Steinsson.** 2018b. "Identification in Macroeconomics." *Journal of Economic Perspectives*, 32(3): 59–86.
- Plagborg-Møller, Mikkel, and Christian K. Wolf.** 2021. "Local Projections and VARs Estimate the Same Impulse Responses." *Econometrica*, 89(2): 955–980.
- Rigobon, Roberto.** 2003. "Identification Through Heteroskedasticity." *The Review of Economics and Statistics*, 85(4): 777–792.
- Sentana, Enrique, and Gabriele Fiorentini.** 2001. "Identification, estimation and testing of conditionally heteroskedastic factor models." *Journal of Econometrics*, 102(2): 143–164.
- Sims, Christopher A.** 2021. "SVAR Identification Through Heteroskedasticity With Misspecified Regimes." Manuscript, Princeton University.
- Stock, James H., and Mark W. Watson.** 2018. "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments." *The Economic Journal*, 128(610): 917–948.