Signing-out Confounding Shocks in Variance-Maximizing Identification Methods*

By NEVILLE FRANCIS AND GENE KINDBERG-HANLON†

Economists have long debated the driver(s) of business cycles, attempting to answer this question through the lens of structural vector autoregressions (SVARs) and estimated DSGE models (Galí, 1999; Smets and Wouters, 2007). More recently, a range of variance-maximizing SVAR estimators have been used to identify some of these potential drivers, including technology shocks (Francis et al., 2014), "news" shocks (Barsky and Sims, 2011), and other attempts to dissect the business-cycle anatomy (Angeletos, Collard and Dellas, 2020). However, identification performance is poor when shocks other than the target of interest also play a nontrivial role in driving volatility at the targeted horizon or frequency, thus confounding the estimation. The result is that these identifications capture a *hybrid* shock rather than a dominant shock (Dieppe, Francis and Kindberg-Hanlon, 2021).

We suggest a simple enhancement to sharpen the variance-maximizing identification procedure that reduces the influence of confounding shocks. This enhancement is to include theoretically-informed sign and magnitude restrictions in the identification stage of the VAR.

When applying our solution of combining sign and magnitude restrictions in the frequency domain, we establish the relevant importance of different classes of shocks in driving the U.S. business cycle. We find that "demand"-type shocks, which drive up inflation and output, explain a roughly similar proportion of business-cycle variation in GDP as "supply"-type drivers of output, which lower inflation.

* The views expressed herein are those of the authors and should not be attributed to the IMF, its Executive Board, or its management.

† Francis: University of North Carolina, Chapel Hill, 3305 Gardener Hall, nfrancis@unc.edu. Kindberg-Hanlon: International Monetary Fund, gkindberg-hanlon@imf.org

I. An example of overlapping shocks in a variance-maximizing identification

Variance-maximizing SVARs identify shocks as those which dominate the variance of a particular variable of interest. However, the objective variance of interest can take several forms; for example, it can reflect the forecast error variance at a specific horizon (Max-Share approach), or it can reflect the variance within a particular frequency domain (Spectral Max-Share approach), reflecting business-cycle or longer-term variance. A reduced-form VAR can be used to compute the objective variance, V (as a function of the variance-covariance matrix of residuals Σ_u and MA coefficient matrix D), modified appropriately to reflect either the forecast error variance at a targeted horizon, k.

$$V = \left(\sum_{\tau=0}^{k-1} D^{\tau} \Sigma_u D^{\tau\prime} \right)$$

Identifying the shock of interest involves the Lagrangian for V:

$$L(\alpha) = \alpha'(V)\alpha - \lambda(\alpha'\alpha - 1)$$

whose first order conditions reduce to solving for the eigenvector associated with the largest eigenvalue of V.

The identified vector α, is then used to generate a single structural shock, $\tilde{A}\alpha$, where $\tilde{A}$ is the Cholesky decomposition of Σ_u. Other structural shocks are left undetermined.

While this approach seeks to identify a dominant structural driver, it is a linear *combination* of structural shocks that often accounts for the largest share of variance. A simple New Keynesian model is used to demonstrate how the variance maximizing methodology can erroneously produce results combining the effects of a demand shock and a supply-side shock, even when the demand shock drives majority of the variance of output:

$$\tilde{y}_t = \frac{-1}{\sigma}\left(i_t - E_t[\pi_{t+1}] - R_t^{NR}\right) + E_t[\tilde{y}_{t+1}] + \eta_t$$

$$\pi_t = \kappa MC_t + \beta E_t[\pi_{t+1}]$$

$$i_t = \phi_y \tilde{y}_t + \phi_\pi \pi_t$$

$$MC_t = (\sigma + \chi) y_t - (1 + \chi)\vartheta_t$$

Where $\tilde{y}_t$ is the output gap, i_t is the nominal interest rate, π_t is inflation, MC_t is marginal costs, and R_t^{NR} is the natural rate of interest. η_t represents a demand (preferences) shock, while ϑ_t reflects a supply-side shock, such as technology.[1]

The solution to the model can be written as:

$$\begin{bmatrix} \tilde{y} \\ \pi \end{bmatrix} = \begin{bmatrix} \Psi_{y\eta} & \Psi_{y\vartheta} \\ \Psi_{\pi\eta} & \Psi_{\pi\vartheta} \end{bmatrix} \begin{bmatrix} \eta_t \\ \vartheta_t \end{bmatrix}$$

Take the highly-simplified example in which both η and ϑ have unit variance and in which the empiricist is searching for the shock which maximizes the initial impact variance of the output gap ($\tilde{y}$). In the case of a parameterization where $\Psi_{y\eta} > \Psi_{y\vartheta}$, the demand shock η drives the largest share of the variance of the output gap. However, an eigenvalue-eigenvector decomposition shows that the variance of $\tilde{y}$ on impact is actually maximized by a combination of η and ϑ (Online Appendix Section 1). More specifically, the "dominant" shock's impact on $\tilde{y}$ will be $\sqrt{\Psi_{y\eta}^2 + \Psi_{y\vartheta}^2}$, while the shock's impact on π will be $(\Psi_{y\eta}\Psi_{\pi\eta} + \Psi_{y\vartheta}\Psi_{\pi\vartheta})/(\sqrt{\Psi_{y\eta}^2 + \Psi_{y\vartheta}^2})$, rather than the true impacts of $\Psi_{y\eta}$ and $\Psi_{\pi\eta}$, respectively.

Note that this identification is equivalent to the standard Cholesky identification solution in this basic case, although maximization over longer periods will deviate from this solution. The smaller supply shock ϑ may exert considerable influence on the properties of the identified shock. Notice for example, that even in cases where the impact of ϑ on the output gap, $\Psi_{y\vartheta}$, is small, the bias to the inflation impulse response function (hereafter, IRF) can still be large if the supply shock's effect on inflation, $\Psi_{\pi\vartheta}$, is large.[2]

In general, the researcher will not restrict her search for the dominant driver of the initial impact variance of the endogenous variables, but rather the forecast error variance at longer horizons, or the variance within a particular frequency band. However, we argue that the same principles shown above still apply; the identified shock will capture a range of influences in proportion to their impacts at the chosen horizon or frequency band.

In summary, without further identifying restrictions, the search for a dominant driver of a variable of interest will be confounded by other shocks.

II. Methodology

To sharpen identification, we propose a maximization procedure that imposes additional restrictions to reduce the influence of shocks that are of less interest to the researcher. Our estimation procedure is to maximize,

$$V(\alpha) = \alpha' V \alpha$$

subject to

$$\alpha'\alpha = 1$$

$$C_R^{L\prime}\alpha \geq a$$

$$\frac{C_R^{NL1\prime}\alpha}{C_R^{NL2\prime}\alpha} \geq b$$

Here, α is chosen as a linear combination of the reduced-form innovations to the variance-covariance matrix of the target variable of interest (V). It also satisfies the unit-length constraint, and is subject to the restriction that it satisfies a set of linear inequality restrictions (C_R^L) and nonlinear inequality restrictions that can be used to regulate the magnitude of the response of variables relative to one another (C_R^{NL}). With the additional inequality constraints, the problem is solved using a constrained maximization algorithm.[3] Our approach differs from standard

[1] σ is the inter-temporal elasticity of substitution, χ is the Frisch elasticity of labor supply. κ is the slope of the Phillips curve and is a function of the probability of not being able to reset prices each period (θ) and the discount rate (β): $\kappa = (1-\theta)(1-\beta\theta)/\theta$.

[2] Our concept of "confounding" shocks in variance-maximizing restrictions has many parallels with the issue of "masquerading" shocks that can lead to misleading results when applying pure sign restrictions (Wolf, 2020).

[3] Additional iterative procedures have been identified which solve constrained eigenvector-eigenvalue decomposition where the linear constraints hold with equality at a. However, even in

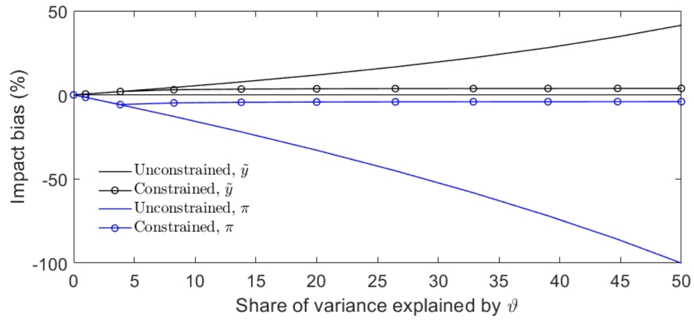

Figure 1. : **Bias of the output gap and inflation impact response to the identified dominant driver of the output gap, with and without restrictions**

sign-restricted identifications in that the draws that satisfy the sign restrictions are not from a uniformly random distribution (Haar prior). Instead, the draws that are kept satisfy the sign restriction constraints and dominate the objective variance function of interest.

III. Constrained Maximization: Applied to the Simple New Keynesian Model

Taking the above New Keynesian model as an example, it is possible to compare the IRF bias that would result from an unconstrained variance maximization identification with a constrained maximization procedure that imposed sign and relative magnitude restrictions. The unconstrained maximization procedure used to capture the dominant driver of the output gap, η, is increasingly biased for both the impact on $\tilde{y}$ and π as the standard deviation of ϑ increases (Figure 1). This bias could lead to the erroneous conclusion that the slope of the Phillips curve is flat, or even non-existent, in response to the main business-cycle driver of the model, a key finding of (Angeletos, Collard and Dellas, 2020).

If the researcher instead imposes *a priori* knowledge of the Philips curve relationship then the bias is substantially lowered and largely invariant to the size of interference from ϑ. The restriction imposed is that inflation increases at least one-third (i.e., $b = \frac{1}{3}$) as much as the increase of the output gap, consistent with standard model parameters. Applying relative magnitude restrictions that are too low or too high are also found to reduce IRF biases for a wide range of tolerances.[4] In addition, the application of relative impact magnitude restrictions is also found to sharpen identification in larger and more complex models and when the objective variance is expressed in frequency-domain form (Online Appendix Section 2).

this case a search algorithm is employed and the solution may have multiple roots (Gander, Golub and von Matt, 1989).

IV. Constrained maximization: What drives the U.S. business cycle?

We now apply this methodology to identify the dominant driver of the variance of U.S. GDP at business-cycle frequencies. V now takes a more complicated form based on a transformation of the MA-coefficient matrix D to capture business-cycle frequencies (ω):

$$V = \left(\sum_{\tau=0}^{k-1} D^{\tau}(e^{-i\tau\omega}) \Sigma_u D^{\tau}(e^{i\tau\omega})' \right)$$

Here, k is set to 40 such that D admits a long-term, but finite series with which to assess the spectral density of the endogenous variables (Dieppe, Francis and Kindberg-Hanlon, 2021).

We estimate a quarterly VAR over the period 1953-2018 containing: log real GDP levels per capita, the cumulative utilization-adjusted TFP log difference series of Fernald (2014), log total hours worked per capita, the unemployment rate, the share of investment in GDP, the share of consumption in GDP, inflation measured by

[4] Applying relative magnitude restraints that are too low perform at least as well as the unrestricted case. Applying elasticity restrictions of up to 0.5 reduce bias in cases where φ explains more than 5 percent of the variance of the output gap for inflation, and reduce the bias of the output gap impact in cases where ϑ explains at least 25 percent of the variance of the output gap.

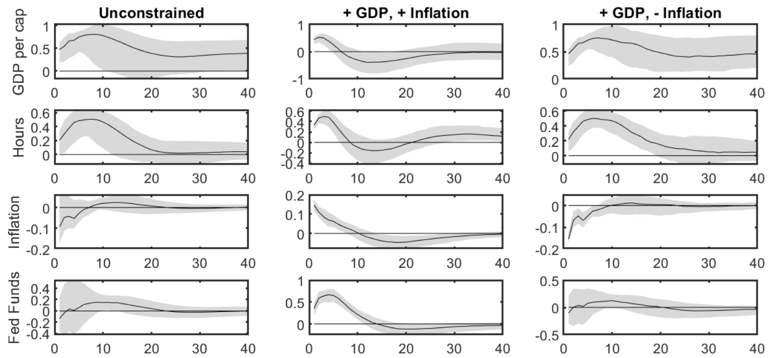

Figure 2. : **Targeting output at business-cycle frequencies, constrained and unconstrained: U.S Data**

Note: 16th and 84th percentile error bands. Columns reflect the unconstrained shock which maximizes business-cycle frequency variation of GDP; the shock where the impact on GDP is constrained to be positive, and the impact on inflation is at least 0.3 times the GDP impact; and, the shock where the impact is constrained to be positive for GDP, but at least -0.3 times the GDP impact for inflation.

the consumption deflator, and the Federal Funds rate of interest (Online Appendix Section 3). The shock which maximizes GDP variation at business-cycle frequencies (6-32 quarters) is identified, in some cases subject to constraints on the response of variables.

The results vary significantly when constraints on the response of inflation are introduced (Figure 2). In the case of a constraint that the initial impact on inflation is at least one-third of the size of the impact on GDP combined with a requirement that the GDP response is positive, there is a less persistent response of GDP relative to the unconstrained case. Furthermore, the response of interest rates is also positive in the demand case, while TFP falls, in contrast to the persistent rise in the unconstrained case (Online Appendix Section 3). In the case of a constraint that inflation *falls* by at least one-third of the increase in GDP, the identified shock takes on the properties of a positive supply-side innovation: the GDP response is still positive at the 10-year horizon, as is the response of TFP and consumption, while interest rates rise very little.

The unconstrained identification produces a hybrid of these two restricted identifications. Even the restricted identifications may continue to be subject to interference, and the imposition of the one-third restriction on the reaction of the inflation rate relative to output is still subject to much debate in the literature. Nonetheless, we argue that the restriction can be varied and still be informative about the contributions of different subsets of shocks to the business cycle.

V. How important are different drivers of the business cycle?

Natural questions that arise from the restricted maximizing shocks are: what proportion of business-cycle variation in output do the newly identified shocks explain relative to the unconstrained case? And, how sensitive are these identified shocks to the relative magnitude restriction on the response of inflation relative to output?

Both restricted-maximization shocks explain around half of the business-cycle variance of GDP. That suggests that both classes of shocks are broadly similar in importance in driving the business cycle (Table 1). As the elasticity restriction is increased, the share of explained variance gradually falls; in the case of both the restrictions requiring a positive and negative response of inflation, the share falls from 53 to 46 percent as the elasticity is increased from 0.05 to 0.4. Clearly, there is likely to be continued overlap between the shocks contained in either identification; it cannot be the case that two independent shocks explain 50 percent or more of the total business-cycle variation of output. Nonetheless, the relative magnitude restrictions go some way to reducing the degree to which different classes of shocks are included. For example, the unrestricted shock explains about 60 percent of the business cycle variation of output.[5]

[5] Furthermore, we find that supply-side drivers of the business-cycle are similarly important in driving long-run variation for GDP (Online Appendix Section 5).

Table 1—: **Contribution of identified shocks to business-cycle variation of GDP**

Scale of restriction	**0.05**	**0.1**	**0.2**	**0.3**	**0.4**
Positive response of inflation	53	53	50	48	46
	(47, 60)	(46, 60)	(43, 57)	(41, 56)	(38, 55)
Negative response of inflation	51	50	49	48	46
	(44, 59)	(43, 59)	(41, 58)	(40, 57)	(39, 55)

Note: The median percent contribution of the SVAR-identified shock to business-cycle frequency variation in GDP as the scale of the inflation response to the identified shock relative to the GDP response is altered. 16th and 84th percentiles shown in brackets.

VI. Conclusion

This paper has highlighted potential shortfalls of employing variance-maximizing SVAR identifications to identify dominant structural drivers. It shows that the identified shock will likely be a composite of shocks that can contain very different properties. Even in cases where a single shock dominates the variance of the target variable of interest, the impulse responses for other variables can be significantly biased. We propose additional restrictions that can be employed to sharpen variance-maximizing identifications. Sign and relative magnitude restrictions are shown in examples of model-generated data to generate IRFs that are closer to those of the true dominant structural shocks. In addition, they can also be used to establish the properties of different categories of shock, for example, those with "demand"-type properties and those resembling supply shocks. However, these restrictions rely on *a priori* knowledge of the structure of the economy. When applied to a VAR estimated on U.S. data, demand shocks, which raise output and inflation, and supply shocks, which raise output but lower inflation, account for a similar proportion of the variance of GDP at business-cycle frequencies.

REFERENCES

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. 2020. "Business-Cycle Anatomy." *American Economic Review*, October 2020(10): 3030–70.

Barsky, Robert B., and Eric R. Sims. 2011. "News shocks and business cycles." *Journal of Monetary Economics*, 58: 273–289.

Dieppe, Alistair, Neville Francis, and Gene Kindberg-Hanlon. 2021. "The Identification of Dominant Macroeconomic Drivers: Coping with Confounding Shocks." ECB Working Paper 2534, Frankfurt, Germany:European Central Bank.

Fernald, John. 2014. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." FRBSF Working Paper, Federal Reserve Bank of San Francisco.

Francis, Neville, Michael T Owyang, Jennifer E Roush, and Riccardo DiCecio. 2014. "A Flexible Finite-Horizon Alternative to Long-Run Restrictions with an Application to Technology Shocks." *Review of Economics and Statistics*, 96(4): 638–647.

Galí, Jordi. 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, 89(1): 249–271.

Gander, Walter, Gene Golub, and Urs von Matt. 1989. "A Constrained Eigenvalue Problem." *Linear Algebra and its Applications*, 114(115): 815–836.

Smets, F., and R. Wouters. 2007. "Shocks and frictions in US business cycles: A Bayesian DSGE approach." *American Economic Review*, 97(3): 586–606.

Wolf, Christian. 2020. "SVAR (Mis)identification and the Real Effects of Monetary Policy Shocks." *American Economic Journal: Macroeconomics*, 12(4): 1–32.