*Macroeconomic Dynamics* (2026), **30**, e28, pp. 1–22
doi:10.1017/S1365100526101023

CAMBRIDGE UNIVERSITY PRESS

ARTICLE

# Variance-Maximizing Identification and the Plague of Confounding Shocks

Alistair Dieppe[1], Neville Francis[2], and Gene Kindberg-Hanlon[3]

[1]European Central Bank, Germany
[2]UNC-Chapel Hill, USA
[3]International Monetary Fund, USA
**Corresponding author:** Neville Francis; Email: nfrancis@unc.edu

**Abstract**

This paper examines the vulnerability of variance-maximizing identification, specifically the Max-Share estimator, to biases arising from confounding shocks. We show analytically that Max-Share tends to recover a hybrid of structural disturbances rather than the dominant long-run shock of interest, even when the target shock clearly drives low-frequency variation. To address this weakness, we propose a band-limited frequency-domain estimator – the Limited-Spectral approach – which isolates low-frequency dynamics and reduces contamination from higher-frequency shocks. Using a transparent two-shock model, we characterize the confounding mechanism and document its implications through Monte Carlo experiments. We further show that, while Limited-Spectral improves identification of low-frequency shocks, variance maximization in the spectral domain can still face difficulties at business-cycle frequencies when multiple structural shocks overlap.

**Keywords:** Identification; long-horizon and business-cycle shocks; confounding shocks

**JEL classifications:** C11; C30; E32

## 1. Introduction

There has been growing use of the Max-Share approach to identify dominant structural shocks. This paper shows that Max-Share suffers from a fundamental weakness that has received little attention: regardless of the targeted horizon, Max-Share often recovers a *mixture* of shocks rather than the single shock of interest. The estimator struggles to disentangle the target disturbance from other innovations in the data-generating process (DGP). We refer to these non-target influences as *confounding shocks*.

Max-Share's appeal stems largely from its ability to sidestep the small-sample difficulties that affect long-run restricted VARs (Francis et al. 2014). Yet the literature has largely overlooked whether variance-maximizing estimators can accurately isolate long-run shocks when confounding disturbances matter. A notable exception is Chari et al. (2008), who demonstrate – within a DSGE framework – that long-run restricted VARs are unbiased only when technology shocks dominate output movements and non-technology shocks contribute minimally to output volatility. When non-technology disturbances play a sizable role, long-run identification becomes biased and yields confidence intervals that are "confidently wrong." Despite evidence that technology shocks may account for anywhere from 2% to 63% of output variability (Galí and Rabanal, 2005; Christiano et al. 2003; Chari et al. 2008), the implications of such confounding for Max-Share and related variance-maximizing estimators remain unexplored.

Recent work has extended variance maximization to the frequency domain, targeting medium- and high-frequency dynamics instead of long-run variation (Angeletos et al. 2020; Giannone et al.

© The Author(s), 2026. Published by Cambridge University Press.

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

2019). For example, Angeletos et al. (2020) identify a dominant business-cycle shock that drives real economic fluctuations with minimal inflationary impact – an interpretation that departs sharply from standard New Keynesian models.

In this paper, we argue that frequency-domain variance maximization should also be applied to low-frequency identification problems. We propose a simple extension of Max-Share – what we call the *Limited-Spectral* estimator – that replaces the long-run variance-covariance matrix with a band-limited measure obtainable reliably in small samples. By restricting attention to the low-frequency region where long-run shocks operate, this approach sharpens Max-Share and often reduces confounding.

Our contribution is complementary to that of Dieppe et al. (2021) and Francis and Kindberg-Hanlon (2022). Dieppe et al. (2021) show that variance-maximizing methods in both time and frequency domains tend to produce hybrid shocks when multiple disturbances are active within the targeted frequency band. Their focus is on interpreting this hybridization, not on mitigating it. Francis and Kindberg-Hanlon (2022) propose adding sign or relative-magnitude restrictions to reduce confounding at business-cycle frequencies. By contrast, we study a different identification problem – the recovery of *long-run* technology shocks – and show that confounding arises even when the target shock dominates very-low-frequency variation. We provide a transparent two-shock model that isolates the mechanism analytically and quantify the resulting bias using Monte Carlo evidence. Most importantly, we offer a tractable solution: a band-limited variance-maximizing estimator that targets only the low-frequency band relevant for long-run technology shocks, reducing contamination without imposing auxiliary restrictions.

Our results speak directly to recent findings by Dou et al. (2025), who document that Max-Share estimators can recover hybrid shocks even when the target shock is the dominant low-frequency driver. Their results reveal the empirical symptom; our analysis identifies the underlying mechanism and proposes an estimator that attenuates it by filtering out the frequencies at which competing shocks operate.

While the Limited-Spectral approach is well suited for isolating long-run technology shocks, it does not provide blanket protection against confounding in all settings. As an exercise, we demonstrate analytically – and using a calibrated New Keynesian model – that variance-maximizing methods applied at business-cycle frequencies can still be distorted by confounding shocks. We replicate the business-cycle findings of Angeletos et al. (2020) using U.S. data and then provide a counterexample using simulated data, showing that spectral variance maximization can misidentify the anatomy of business-cycle fluctuations.

A related strand of empirical work argues that variance–maximizing identification can, under specific conditions, recover economically meaningful shocks. In particular, Ben Zeev (2018, 2019) show that when identification incorporates an equipment-based relative price of investment and restricts its long-run behavior to two investment-specific technology (IST) shocks, the variance–maximizing shock behaves empirically like IST news. This result suggests that, in environments where the equipment price exhibits a distinctive long-run structure closely tied to IST improvements, Max-Share may confirm the dominant role of IST news in driving business-cycle fluctuations.

We revisit this mechanism in our empirical analysis. Using an equipment-based relative price of investment series and imposing the same agnostic two-shock long-run restriction, we find that the resulting RPI-anchored Max-Share shock does not generally align with IST news. Moreover, the unrestricted and Limited-Spectral Max-Share shocks remain nearly identical and differ substantially from IST news. These findings indicate that the empirical equivalence between Max-Share and IST news documented in earlier work is not generic but instead depends on the presence of a sufficiently distinctive long-run pattern in the data. Absent such structure, variance–maximizing identification continues to produce hybrid shocks, as predicted by our theoretical analysis.

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

The remainder of the paper proceeds as follows. Section 2 reviews the Max-Share approach. Section 3 demonstrates its susceptibility to confounding shocks and introduces our Limited-Spectral estimator. Section 4 develops theoretical intuition and presents Monte Carlo evidence using a simple two-shock model. Section 5 applies our framework to business-cycle frequencies, illustrating how spectral variance–maximization may also fail in that setting. Section 6 examines empirical conditions under which Max-Share may or may not align with IST news shocks. Section 7 concludes.

## 2. Variance-maximizing approach to identification

In this section, we provide a brief summary of the variance-maximizing approach (of which Max-Share is a variant) for identifying *dominant* structural shocks. Variance-maximizing identification proceeds by solving an eigenvalue–eigenvector problem associated with a quadratic form in the forecast error variance. [1] Specifically, let $V_\tau$ denote the symmetric, positive semi-definite matrix governing the contribution of candidate shocks to the forecast error variance at horizon $\tau$, and let $\alpha$ be a unit-length vector. The problem can be written as

$$\max_{\alpha:\,\alpha'\alpha=1} \alpha' V_\tau \alpha.$$

Setting up the corresponding Lagrangian,

$$L(\alpha) = \alpha' V_\tau \alpha - \lambda(\alpha'\alpha - 1),$$

the first-order conditions imply that the solution $\alpha$ is the eigenvector associated with the largest eigenvalue of $V_\tau$,

$$V_\tau \alpha = \lambda \alpha.$$

### *2.1 Max-Share*

The Max-Share approach to identification assumes that a particular shock is the *dominant* driver over a specific horizon. The shock of interest is that which drives the largest proportion of the forecast error variance at the chosen horizon (see Francis et al. (2014)). For long-run applications, Francis et al. (2014) surmises that 10 years is long enough to capture dominant long-run shocks.

The forecast error at horizon $k$ can be written:

$$y_{t+k} - \hat{y}_{t+k} = \sum_{\tau=0}^{k-1} D_\tau u_{t+k-\tau}, \tag{1}$$

where $D_\tau$ is the matrix of moving average coefficients. By defining $\tilde{A}_0$ as the Cholesky decomposition of the variance-covariance matrix $\Sigma_u$ and $\varphi_i$ as a selection vector (size $n \times 1$) with 1 in the $i$-th position, we find the shock vector $\alpha$ which maximizes the contribution to the total forecast error variance of variable $i$ at horizon $k$:

$$\max_{\alpha} f(\alpha) = \frac{\varphi_i' \left( \sum_{\tau=0}^{k-1} D_\tau \tilde{A}_0 \alpha (\tilde{A}_0 \alpha)' D_\tau' \right) \varphi_i}{\varphi_i' \left( \sum_{\tau=0}^{k-1} D_\tau \Sigma_u D_\tau' \right) \varphi_i} \quad \text{s.t.} \quad \alpha'\alpha = 1. \tag{2}$$

The identified structural shock is defined by the impact vector $\tilde{A}_0\alpha$. Following Uhlig (2003), the vector $\alpha$ that solves this maximization problem is the eigenvector associated with the maximum eigenvalue of the matrix $S_{i,k}$. This matrix is derived by rearranging the numerator in (2) to isolate $\alpha$:

$$S_{i,k} = \sum_{\tau=0}^{k-1} \tilde{A}_0' D_\tau' \varphi_i \varphi_i' D_\tau \tilde{A}_0. \tag{3}$$

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

Note that the denominator in (2) represents the *total* forecast error variance of variable $i$, which is a scalar constant independent of $\alpha$ and therefore does not affect the first-order conditions of the maximization.

#### *2.1.1 First warning on the use of Max-Share*

When employing Max-Share, using more lags of the reduced-form MA coefficients does not necessarily translate into a *better* estimate. Researchers have to be cognizant of the fact that as the horizon lengthens there are (cumulative) errors in the MA coefficients on which the identification rests. These errors can be sizeable as we argue in Lemma 1:

**Lemma 1.** *VARs can produce increasingly biased and uncertain impulse responses as the horizon increases. A simple argument illustrates this issue.*

*Consider the AR(1) model*

$$y_t = \phi y_{t-1} + \varepsilon_t, \qquad |\phi| < 1.$$

*Let $\hat{\phi}$ be an estimator of $\phi$ such that $\hat{\phi} \xrightarrow{p} \phi$ and*

$$\sqrt{T}\,(\hat{\phi} - \phi) \xrightarrow{d} \mathcal{N}\left(0, \sigma_\phi^2\right).$$

*Define the j-step impulse response $\Psi_j = \phi^j$ and its plug-in estimator $\hat{\Psi}_j = \hat{\phi}^j$ for a fixed integer $j \geq 1$. Then, by the delta method,*

$$\sqrt{T}\left(\hat{\Psi}_j - \Psi_j\right) \xrightarrow{d} \mathcal{N}\left(0,\ \left(j\phi^{j-1}\right)^2 \sigma_\phi^2\right).$$

*Moreover, a second-order delta-method expansion yields the approximation*

$$\mathbb{E}[\hat{\Psi}_j] \approx \phi^j + \frac{1}{2}\, j(j-1)\, \phi^{j-2}\, Var(\hat{\phi}),$$

*and therefore the (approximate) bias satisfies, for $j \geq 2$,*

$$\mathbb{E}[\hat{\Psi}_j] - \Psi_j \approx \frac{1}{2}\, j(j-1)\, \phi^{j-2}\, Var(\hat{\phi}).$$

*If $\hat{\phi}$ is additionally biased with $\mathbb{E}[\hat{\phi}] = \phi + b$ and $Var(\hat{\phi}) = v$ (with $b$, $v$ small), then*

$$\mathbb{E}[\hat{\Psi}_j] - \Psi_j \approx j\phi^{j-1} b + \frac{1}{2}\, j(j-1)\, \phi^{j-2} v.$$

*In particular, even if $b < 0$, the impulse response estimator can be upward biased when*

$$\frac{1}{2}(j-1)v > -\phi\, b.$$

## 3. Unpacking confounding shocks

### *3.1 Demonstrating the Impacts of Confounding Shocks: Second Warning to the Use of Max-Share*

The prevailing approach in the literature is to use Max-Share to identify long-run shocks. For a simple two-variable VAR consisting of labor productivity ($p$) and hours, ($n$), with the true structural form:

$$A_0 \epsilon = \begin{bmatrix} A_{lp} & A_{hp} \\ A_{ln} & A_{hn} \end{bmatrix} \begin{bmatrix} \epsilon^l \\ \epsilon^h \end{bmatrix},$$

where we use $l$ and $h$ to characterize the first shock as a low-frequency (technology shock) and the second as a high-frequency (business cycle) shock. In the simplest case, at horizon $k = 0$,

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

the solution's impact on $p$ and $n$ will match the impacts using a solution based on the Cholesky decomposition of $\Sigma_u$.

$$\tilde{A}_0\tilde{q}(0) = \begin{bmatrix} \sqrt{A_{lp}^2 + A_{hp}^2} \\ \frac{A_{lp}A_{ln}+A_{hp}A_{hn}}{\sqrt{A_{lp}^2+A_{hp}^2}} \end{bmatrix}$$

The impact of the identified shock on productivity will be a combination of the low- and high-frequency structural shocks, in proportion to their respective impacts. The impact on hours will depend on the relative impact on productivity and the relative impact on hours of each shock.

Whilst intuitive, the above expression is an oversimplification of the effects of confounding shocks. This is because we have only demonstrated its effects on the impact impulse responses. Often, however, the practitioner is interested in longer-run applications and thus wants to identify low-frequency shocks operating at long(er) horizons. In these cases, the solution becomes even more complex (Appendix A). Nevertheless, the identified dominant shock at longer horizons will be a hybrid shock, capturing elements of the underlying structural drivers in proportion to their initial impacts and persistence.

### *3.2 Establishing a More Robust Estimator*

To reduce the influence of confounding shocks, we need to be able to filter out these competing shocks from the maximization to the greatest extent possible so that the *true shock* is as dominant as possible at the horizon of interest. To this end, we appeal to a truncated-kernel spectral density. We interpret the difference between our (time-domain) Max-Share approach and its truncated-kernel spectral counterpart as one plausible measure of the extent to which confounding shocks impact Max-Share. However, we italicize *more robust* in anticipation of a future cautionary caveat.

#### *3.2.1 Spectral identification*

The spectral approach to identifying dominant structural shocks operates in the frequency domain by maximizing the share of variance of a specific target variable explained within a prespecified range of frequencies.

To implement this methodology, consider the Wold representation

$$Y_t = \sum_{j=0}^{\infty} D_j u_{t-j}, \qquad \mathbb{E}[u_t u_t'] = \Sigma_u.$$

and define the autocovariance matrices

$$\Gamma(\tau) = \mathbb{E}[Y_t Y_{t-\tau}'], \qquad \tau \in \mathbb{Z}.$$

The spectral density matrix is given by

$$S_{YY}(\omega) = \frac{1}{2\pi} D\left(e^{-i\omega}\right) \Sigma_u D\left(e^{i\omega}\right)', \qquad D\left(e^{-i\omega}\right) = \sum_{j=0}^{\infty} D_j e^{-i\omega j}, \tag{4}$$

and satisfies the Fourier relationship

$$\Gamma(\tau) = \int_{-\pi}^{\pi} e^{i\omega\tau} S_{YY}(\omega)\, d\omega, \qquad S_{YY}(\omega) = \frac{1}{2\pi} \sum_{\tau=-\infty}^{\infty} \Gamma(\tau) e^{-i\omega\tau}. \tag{5}$$

To identify the shock, we define $\tilde{A}_0$ as the Cholesky decomposition of $\Sigma_u$ and $\varphi_i$ as a selection vector for the target variable $i$. We seek the structural shock vector $\alpha$ (where $\alpha'\alpha = 1$) that

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

maximizes the variance of variable $i$ within the frequency band $[\underline{\omega}, \bar{\omega}]$:

$$\max_{\alpha} f(\alpha) = \frac{\varphi_i' \left( \int_{\underline{\omega}}^{\bar{\omega}} D(e^{-i\omega}) \tilde{A}_0 \alpha (\tilde{A}_0' \alpha) D(e^{i\omega})' \, d\omega \right) \varphi_i}{\varphi_i' \left( \int_{-\pi}^{\pi} S_{YY}(\omega) \, d\omega \right) \varphi_i}. \tag{6}$$

Analogous to the Max-Share approach, the vector $\alpha$ that solves this maximization corresponds to the eigenvector associated with the largest eigenvalue of the matrix $V_i(\underline{\omega}, \bar{\omega})$, which isolates the contribution of the shock to the target variable:

$$V_i(\underline{\omega}, \bar{\omega}) = \int_{\underline{\omega}}^{\bar{\omega}} \text{Re} \left( \tilde{A}_0' D(e^{i\omega})' \varphi_i \varphi_i' D(e^{-i\omega}) \tilde{A}_0 \right) \, d\omega.$$

#### *3.2.2 Limited-horizon spectral identification*

One criticism that can be leveled against the spectral identification approach is that the long-run VAR representation used to calculate the spectrum of the endogenous variables may be biased when estimated on a short sample of data and with limited lags. A "windowed" selection of autocorrelations is often used to estimate the spectrum of a series. The same principle can be used in the spectral VAR identification process, effectively applying a truncated kernel estimation approach. For example, the infinite-MA representation, $D(e^{-i\tau\omega})$, can be replaced with a limited horizon series of impulse response coefficients.

$$D^k(e^{-i\omega}) = \sum_{j=0}^{k-1} D_j e^{-ij\omega}. \tag{7}$$

The same procedure used to identify the shock of interest in the general spectral approach can be straightforwardly applied here as well. Additionally, applying this truncation provides a direct comparison to Max-Share. That is, both Max-Share and Limited-Spectral methods will use the same horizon $k$ in their identification step, making the only difference between the two approaches the additional filtering of the Limited-Spectral approach at specific frequencies. It is this feature of the Limited-Spectral method to which we appeal when interpreting anything extraneous coming from Max-Share as potentially reflecting confounding shocks that influence the data at frequencies within the excluded domain.

### *3.3 Quantifying Confounding Shocks*

In population and abstracting from spectral estimation error, the Limited-Spectral identification can be interpreted as a frequency-restricted analogue of the Max-Share procedure, emphasizing variation within a prespecified band. In estimation, differences between the two approaches may therefore be informative about confounding dynamics but can also arise from finite-sample spectral estimation error or frequency leakage. Accordingly, we interpret such differences as a diagnostic rather than as a definitive measure of confounding shocks.

The forecast error variance matrix with respect to which the maximization is performed can be written, for the Max-Share identification, as

$$V_k^{\text{Max-Share}} = \sum_{\tau=0}^{k-1} D^\tau \Sigma_u D^{\tau\prime}, \tag{8}$$

and for the Limited-Spectral identification, as

$$V_k^{\text{L-Spectral}} = \int_{\underline{\omega}}^{\bar{\omega}} \text{Re} \left( \sum_{\tau=0}^{k-1} D^\tau \left( e^{i\tau\omega} \right) \, \Sigma_u \, D^\tau \left( e^{-i\tau\omega} \right)' \right) d\omega. \tag{9}$$

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

More formally, Equation 9 corresponds to a finite-sample discrete-frequency approximation to the band-integrated spectral density obtained by replacing the infinite-horizon moving-average representation with its truncated counterpart. This construction is equivalent to applying a windowed spectral estimator that emphasizes variation at the selected frequencies while maintaining comparability with the finite-horizon Max-Share criterion. Differences between the two identified impulse responses therefore reflect the exclusion of variation outside the targeted frequency band and may capture the influence of dynamics operating at excluded frequencies. For example, comparing Max-Share applied at long horizons with the Limited-Spectral identification at low frequencies can be informative about the role of higher-frequency shocks in shaping long-horizon dynamics.

*From this point forward, we focus exclusively on the Limited-Spectral approach in the frequency domain. For brevity, we occasionally refer to this method simply as the* ***spectral approach****.*

## 4. Theoretical simulations: establishing limited-spectral as a robust alternative in low-frequency-dominant environments

This section explains the simulation framework used to study the performance of the Max-Share and Limited-Spectral estimators in a setting that mirrors our empirical application. We first outline the construction of a two-variable model that generates artificial data for productivity and hours. We then describe how impulse responses and identified shocks are recovered from an estimated VAR and how we assess the accuracy of each estimator using Monte Carlo experiments. Finally, we interpret the results through a spectral lens and discuss their implications for the business-cycle application.

### *4.1 Two-Variable Model and Identification*

We work with a simple two-variable DGP in which labor productivity $L_t$ and hours $N_t$ are driven by a technology shock $\epsilon_t^z$ and a non-technology shock $\epsilon_t^b$:

$$L_t = z_t^l + b_t, \tag{10}$$

$$N_t = 0.7N_{t-1} - 0.3N_{t-2} - 0.3z_t^l + 0.3b_t, \tag{11}$$

$$z_t^l = \rho_z z_{t-1}^l + z_t^g, \tag{12}$$

$$z_t^g = \rho_{z^g} z_{t-1}^g + \epsilon_t^z, \tag{13}$$

$$b_t = \rho_{b,1} b_{t-1} + \rho_{b,2} b_{t-2} + \epsilon_t^b. \tag{14}$$

The technology innovation $\epsilon_t^z$ drives a highly persistent level component $z_t^l$ and, through $z_t^g$, a growth component that can itself be persistent. The non-technology disturbance $b_t$ is a stationary but cyclical AR(2) process that acts as a confounding driver of $L_t$ and $N_t$.

In the baseline implementation we simulate $T = 2{,}000$ usable observations and discard an initial burn-in of 500 periods to eliminate the influence of initial conditions. The structural shocks are i.i.d. Gaussian,

$$\epsilon_t^z \sim \mathcal{N}\big(0, \sigma_z^2\big), \qquad \epsilon_t^b \sim \mathcal{N}\big(0, \sigma_b^2\big),$$

with $\sigma_z = 1$ throughout. The remaining parameters $(\rho_z, \rho_{z^g}, \rho_{b,1}, \rho_{b,2}, \sigma_b)$ are calibrated so that the technology component is highly persistent while the non-technology component is less persistent but relatively volatile. In the baseline stationary case we set

$$\rho_z = 0.9, \qquad \rho_{z^g} = 0, \qquad \sigma_z = 1,$$

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

and parameterize the confounding disturbance $b_t$ as a business-cycle AR(2) process with complex roots. Specifically, we choose

$$\rho_{b,1} = 2R\cos(\lambda), \qquad \rho_{b,2} = -R^2, \qquad R = 0.85, \qquad \lambda = 2\pi/8, \qquad \sigma_b = 2.$$

This calibration concentrates the spectral mass of the confounding shock at cyclical frequencies (roughly an 8-quarter cycle), while the technology component remains highly persistent and contributes disproportionately at very low frequencies. The resulting DGP therefore isolates the case of *business-cycle confounding*: the non-technology shock generates substantial medium-frequency variation that can contaminate variance-maximizing identification in finite samples.

Given a realization of $(L_t, N_t)$, we estimate a bivariate VAR of order two,

$$Y_t = A_1 Y_{t-1} + A_2 Y_{t-2} + u_t, \qquad Y_t = \begin{pmatrix} L_t \\ N_t \end{pmatrix},$$

using Gaussian maximum likelihood. Let $\Sigma_u = \mathrm{Var}(u_t)$ denote the reduced-form covariance matrix, and let $A$ be its lower-triangular Cholesky factor, $\Sigma_u = AA'$.

From the estimated VAR we construct the moving-average (MA) representation

$$Y_{t+h} = \sum_{\tau=0}^{\infty} D_\tau u_{t+h-\tau},$$

where $D_0 = I_2$ and, for $h \geq 1$, the MA coefficients satisfy the VAR(2) recursion

$$D_h = A_1 D_{h-1} + A_2 D_{h-2}.$$

We work with a finite horizon $H = 20$ and use $\{D_h\}_{h=0}^{H-1}$ both to form impulse responses and to construct the identification criteria.

**Max-Share identification.** Let $e_1 = (1, 0)'$ select labor productivity. For a given impact direction $a \in \mathbb{R}^2$, the implied response at horizon $h$ is $D_h A a$, and the contribution of the associated shock to the forecast error variance of $L_t$ over horizons $0, \ldots, H-1$ is summarized by

$$a' H^{MS} a, \qquad H^{MS} = \sum_{h=0}^{H-1} (D_h A)' (e_1 e_1') (D_h A).$$

The Max-Share shock is defined as the unit vector $a$ that maximizes this quadratic form. In the implementation this is obtained as the leading eigenvector $\alpha_{MS}$ of the symmetrized matrix $\frac{1}{2}(H^{MS} + (H^{MS})')$, normalized to unit Euclidean length. The corresponding impact vector is $A_0^{MS} = A\alpha_{MS}$ and the Max-Share impulse responses are

$$\mathrm{IRF}_h^{MS} = D_h A_0^{MS}, \quad h = 0, \ldots, H-1.$$

**Limited-spectral identification.** The Limited-Spectral estimator maximizes the contribution of $L_t$ at low frequencies. For a frequency band $[-\omega_B, \omega_B]$ we construct the spectral objective

$$H^{LS} = \frac{1}{J} \sum_{j=1}^{J} \left( D(e^{-i\omega_j})A \right)' (e_1 e_1') \left( D(e^{-i\omega_j})A \right),$$

where $\omega_j$ are $J$ equally spaced grid points in $[-\omega_B, \omega_B]$ and

$$D(e^{-i\omega}) = \sum_{h=0}^{H-1} D_h e^{-i\omega h}$$

is the finite-horizon frequency-domain representation of the MA. In the implementation we use $J = 500$ grid points and set the low-frequency band width to $\omega_B = 2\pi/10$, corresponding to cycles

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

longer than roughly ten periods. As in the Max-Share case, the Limited-Spectral shock is the leading eigenvector $\alpha_{LS}$ of $\frac{1}{2}(H^{LS} + (H^{LS})')$; the associated impact vector is $A_0^{LS} = A\alpha_{LS}$ and the Limited-Spectral IRFs are

$$\mathrm{IRF}_h^{LS} = D_h A_0^{LS}, \quad h = 0, \ldots, H-1.$$

**Spectral intuition.** Both criteria admit a frequency-domain interpretation. Let $S_{YY}(\omega)$ denote the reduced-form spectral density of $Y_t$ and $S_{LL}(\omega)$ its $(1,1)$ element. For a given direction $a$, the contribution of the associated shock to the spectrum of $L_t$ at frequency $\omega$ is proportional to

$$g(\omega, a) = \left(D(e^{-i\omega})Aa\right)' e_1 e_1' \left(D(e^{-i\omega})Aa\right).$$

The Max-Share criterion aggregates $g(\omega, a)$ over the full frequency range, reflecting the fact that the finite-horizon forecast error variance can be written as an integral over $[-\pi, \pi]$. In contrast, the Limited-Spectral criterion integrates $g(\omega, a)$ only over the low-frequency band $[-\omega_B, \omega_B]$. In our calibration the confounding disturbance $b_t$ is concentrated at business-cycle frequencies (around an 8-quarter cycle), while the technology component is highly persistent and therefore contributes disproportionately near zero frequency. As a result, restricting attention to cycles longer than roughly ten quarters shifts the maximizer away from the cyclical confounder and toward the truly low-frequency technology-driven component.

**True IRFs and shock correlations.** Because the DGP is known, we can compute the *true* IRFs of $(L_t, N_t)$ to a one-time technology innovation directly from the structural model. Setting $\epsilon_0^z = 1$ and $\epsilon_t^z = \epsilon_t^b = 0$ for $t \geq 1$, and starting from zero initial conditions, we iterate

$$\begin{aligned}
z_t^g &= \rho_{z^g} z_{t-1}^g + \epsilon_t^z, \\
z_t^l &= \rho_z z_{t-1}^l + z_t^g, \\
b_t &= \rho_{b,1} b_{t-1} + \rho_{b,2} b_{t-2}, \\
L_t &= z_t^l + b_t, \\
N_t &= 0.7 N_{t-1} - 0.3 N_{t-2} - 0.3 z_t^l + 0.3 b_t,
\end{aligned}$$

for $t = 0, \ldots, H-1$. This yields the benchmark IRFs $\mathrm{IRF}_h^{true} = (L_h, N_h)'$.

To evaluate identification performance, we also compare the estimated shocks with the true structural shocks $(\epsilon_t^z, \epsilon_t^b)$. From the estimated VAR we obtain the reduced-form residuals $u_t$ and define the candidate structural shocks

$$\hat{s}_t^{MS} = \alpha_{MS}' A^{-1} u_t, \qquad \hat{s}_t^{LS} = \alpha_{LS}' A^{-1} u_t.$$

Because the sign of an eigenvector is arbitrary, we fix the sign of each estimator so that its IRF for productivity is aligned with the true productivity IRF.[2] We then report correlations with the true shocks in absolute value,

$$\left|\mathrm{corr}\left(\hat{s}_t^{MS}, \epsilon_t^z\right)\right|, \quad \left|\mathrm{corr}\left(\hat{s}_t^{LS}, \epsilon_t^z\right)\right|, \quad \left|\mathrm{corr}\left(\hat{s}_t^{MS}, \epsilon_t^b\right)\right|, \quad \left|\mathrm{corr}\left(\hat{s}_t^{LS}, \epsilon_t^b\right)\right|,$$

as well as pointwise absolute IRF bias, $|\mathrm{IRF}_h^{\cdot} - \mathrm{IRF}_h^{true}|$.

### *4.2 Monte Carlo Evaluation*

We assess the finite-sample performance of the Max-Share and Limited-Spectral estimators through Monte Carlo experiments. For each of $R = 500$ replications we:

1. Draw sequences of $(\epsilon_t^z, \epsilon_t^b)$ and simulate the structural model for $T + 500$ periods using the baseline calibration described above, discarding the first 500 observations.

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

**Table 1.** Monte Carlo correlation of VAR-identified shocks with true shocks

| | Median $\lvert\text{corr}(\hat{s}_t, \epsilon_t)\rvert$ | |
|---|---|---|
| | Technology shock $\epsilon_t^z$ | Non-technology shock $\epsilon_t^b$ |
| Max-Share | 0.578 [0.408, 0.725] | 0.806 [0.658, 0.894] |
| Limited-Spectral | 0.813 [0.697, 0.872] | 0.534 [0.380, 0.698] |

*Notes:* Entries report the median absolute correlation between the estimated shock $\hat{s}_t$ and the true structural shocks $(\epsilon_t^z, \epsilon_t^b)$ across $R = 500$ Monte Carlo replications of the two-variable model. Brackets show the 5th and 95th percentiles of the Monte Carlo distribution.

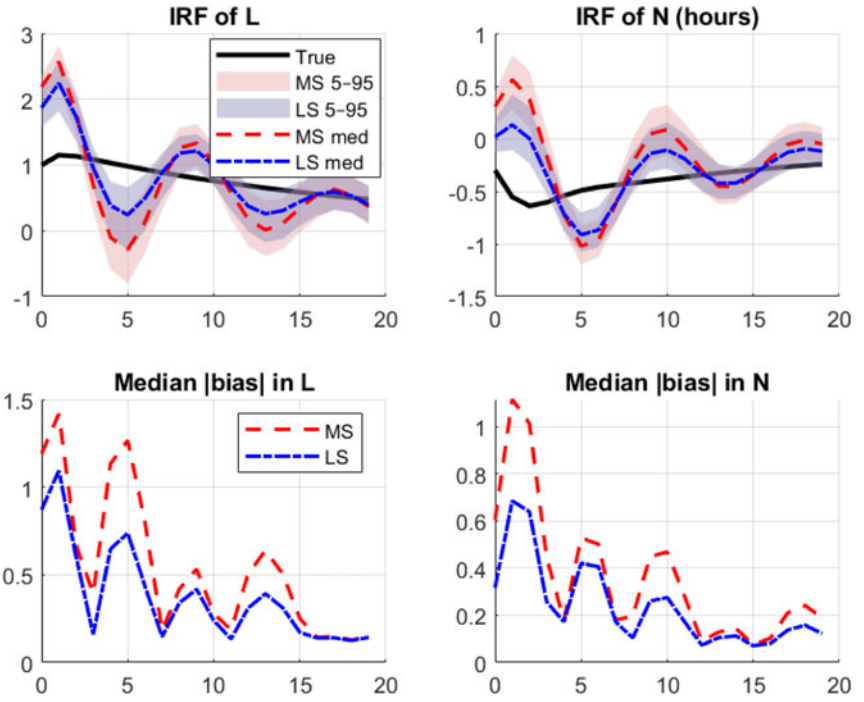



**Figure 1.** Monte Carlo impulse responses and bias in the two-variable model.
*Notes:* Each panel summarizes $R = 500$ Monte Carlo replications. Top-left: IRF of productivity $L$ to a technology shock. Top-right: IRF of hours $N$. Shaded regions denote 5th–95th percentiles across replications. Dashed lines give median IRFs under Max-Share (MS) and Limited-Spectral (LS) identification. Bottom panels show the median absolute deviation from the true IRF for each horizon.

2. Estimate a VAR(2) in levels on the remaining $T = 300$ observations of $(L_t, N_t)$.[3]
3. Construct the MA coefficients $\{D_h\}_{h=0}^{H-1}$, compute the Max-Share and Limited-Spectral impact vectors $A_0^{MS}$ and $A_0^{LS}$, and obtain the corresponding IRFs $\text{IRF}_h^{MS}$ and $\text{IRF}_h^{LS}$.
4. Recover the identified shocks $(\hat{s}_t^{MS}, \hat{s}_t^{LS})$ from the VAR residuals, align their sign with the true productivity IRF, and compute the absolute correlations with $(\epsilon_t^z, \epsilon_t^b)$ as well as the absolute IRF bias relative to $\text{IRF}_h^{true}$.

Figure 1 summarizes the Monte Carlo distribution of IRFs and biases, while Table 1 reports correlations between the estimated shocks and the true structural innovations. In this business-cycle-confounding environment, the distinction between Max-Share and Limited-Spectral is visible both in the recovered shocks and in the induced impulse responses.

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

First, Table 1 shows that Limited-Spectral produces a substantially cleaner estimate of the technology shock. The median absolute correlation with the true technology innovation is

$$\left|\mathrm{corr}\left(\hat{s}_t^{MS}, \epsilon_t^z\right)\right|_{\text{median}} = 0.578 \ [0.408, 0.725], \qquad \left|\mathrm{corr}\left(\hat{s}_t^{LS}, \epsilon_t^z\right)\right|_{\text{median}} = 0.813 \ [0.697, 0.872].$$

Conversely, Max-Share loads much more strongly on the confounding non-technology shock:

$$\left|\mathrm{corr}\left(\hat{s}_t^{MS}, \epsilon_t^b\right)\right|_{\text{median}} = 0.806 \ [0.658, 0.894], \qquad \left|\mathrm{corr}\left(\hat{s}_t^{LS}, \epsilon_t^b\right)\right|_{\text{median}} = 0.534 \ [0.380, 0.698].$$

These statistics align directly with the design of the experiment: because $b_t$ is concentrated at business-cycle frequencies, unrestricted variance maximization can inadvertently select directions that capture substantial cyclical variance even when the object of interest is the persistent technology component. The Limited-Spectral restriction attenuates this mechanism by excluding variance outside the low-frequency band.

Second, Figure 1 shows how this contamination translates into impulse-response distortions: even though technology dominates asymptotically at very low frequencies, finite-horizon variance maximization remains sensitive to cyclical components when identification is performed using truncated impulse responses. The top panels plot the true technology IRFs (black) against the Monte Carlo distribution of identified IRFs under Max-Share (red) and Limited-Spectral (blue). The estimated responses exhibit pronounced cyclical dynamics – reflecting the business-cycle component embedded in the confounding shock – and therefore deviate from the monotone, highly persistent technology-driven path implied by the DGP. The bottom panels quantify these deviations: the median absolute IRF bias is largest at short horizons and declines with the horizon as the transitory cyclical component dissipates, with Limited-Spectral systematically, particularly at short horizons, reducing the bias relative to Max-Share.

Overall, the Monte Carlo evidence illustrates the central message of this section in its cleanest form. When confounding disturbances have substantial power at cyclical frequencies, time-domain variance maximization tends to recover a hybrid shock that is highly correlated with the confounder and generates biased IRFs. Imposing a low-frequency restriction in the objective function materially improves identification of the technology innovation and reduces IRF bias, though it does not eliminate bias entirely in finite samples. This improvement relies on spectral separation between the target and confounding shocks; when multiple disturbances overlap within the low-frequency band itself, frequency restriction need not alter the variance-maximizing solution – a case we document empirically below.

### *4.3 Interpretation and Link to the Business-Cycle Application*

The two-variable simulations are deliberately stylized, but they capture a key feature of the environment where Limited-Spectral is most informative: one component is highly persistent and dominates very low-frequency movements in productivity, while confounding shocks are concentrated at higher (here, business-cycle) frequencies. By construction, the Limited-Spectral objective assigns weight only to a narrow low-frequency band in which the technology shock has relatively greater power, whereas the Max-Share criterion aggregates contributions from all frequencies and is therefore more exposed to variance generated by cyclical confounding disturbances.

The situation will be more challenging in the business-cycle application below, where multiple structural shocks have substantial spectral mass in the same range of frequencies. When both demand and supply shocks contribute appreciably to the variance of output and inflation at business-cycle frequencies (for example, 6–32 quarters – in our empirical applications, we map calendar frequencies into angular frequencies as $\omega = 2\pi/T$, so that business-cycle fluctuations over 6–32 quarters correspond to $\omega \in [2\pi/32,\ 2\pi/6]$), neither time- nor frequency-domain variance-maximizing identification can cleanly isolate a single structural innovation. In that case, both Max-Share and Limited-Spectral procedures will recover hybrids whose interpretation requires additional economic restrictions or sign constraints.

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

The two-variable simulations thus play a dual role. They provide a transparent benchmark in which the gains from Limited-Spectral identification are easiest to see, and they highlight the conditions under which such gains are likely to be realized in practice: a dominant low-frequency driver and confounding shocks that matter relatively more at shorter horizons. In the empirical sections that follow, we use these insights to guide the choice of frequency band and horizon and to interpret the resulting shocks in environments where low-frequency dominance may fail or where the shocks of interest are inherently multi-frequency.

## 5. Caveat to the spectral approach: looking at business-cycle drivers

The robustness of Limited-Spectral hinges on spectral separation; we now examine environments where this condition fails. We have treated the Limited-Spectral method as more immune (or close to immune) to confounding shocks in the identification of technology shocks and, as such, better able to accurately uncover the true impulse responses. However, we now highlight, both analytically and in simulation, a scenario where the spectral approach has difficulty. This example is at business cycle frequency, where there are many (more) drivers and isolating the dominant one is more of a challenge.

### *5.1 Confounding the Identification of Business-Cycle Frequency Shocks*

Recently, some influential papers have employed the spectral approach to Max-Share, as presented here, to identify the driving forces of business cycles. We note that the same logic of confounding shocks could also apply in the case of overlapping business cycle frequency shocks. For example, Giannone et al. (2019) find a shock that maximizes the variance of a target variable at business-cycle frequencies will encompass a linear combination of structural shocks. We formalize this statement, finding that the shock will encompass a combination of structural shocks broadly in proportion to their relative contribution to the variance over the select frequency band. We propose that the same logic outlined in Appendix A also holds in the case of targeting the shock that maximizes the variance within a business-cycle frequency band.

Take, for example, two structural business-cycle frequency shocks, $\epsilon_{BC1}$ and $\epsilon_{BC2}$, affecting the (same) two endogenous variables in the previous example, but now at business-cycle frequencies. As in the example for low-frequency confounding shocks, the identified shock will be a combination of the shocks. In this case, those that drive business cycle frequency variation of the target variable.

This will make it difficult to interpret IRFs of the identified shock in question, given they will represent a combination of drivers with only an *a priori* understanding of their relative weights. Furthermore, the variance of many macroeconomic drivers in the business-cycle frequency domain are more prone to contamination than in the low-frequency range. For example, at business-cycle frequencies, a wider array of shocks influence real and nominal variables. The competing effects of supply-side and demand-side influences may therefore be difficult to disentangle.

### *5.2 Demonstrating the Potential for Confounding Shocks Affecting the Spectral Approach*

In this section, we use an example from the business cycle anatomy literature to demonstrate the potential for confounding shocks to pervert the results at higher frequencies. We first replicate the empirical findings of Angeletos et al. (2020). We then proceed to use a simple theoretical construct to show the possibility that confounding shocks could distort conclusions about the main macroeconomic drivers at business-cycle frequencies.

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

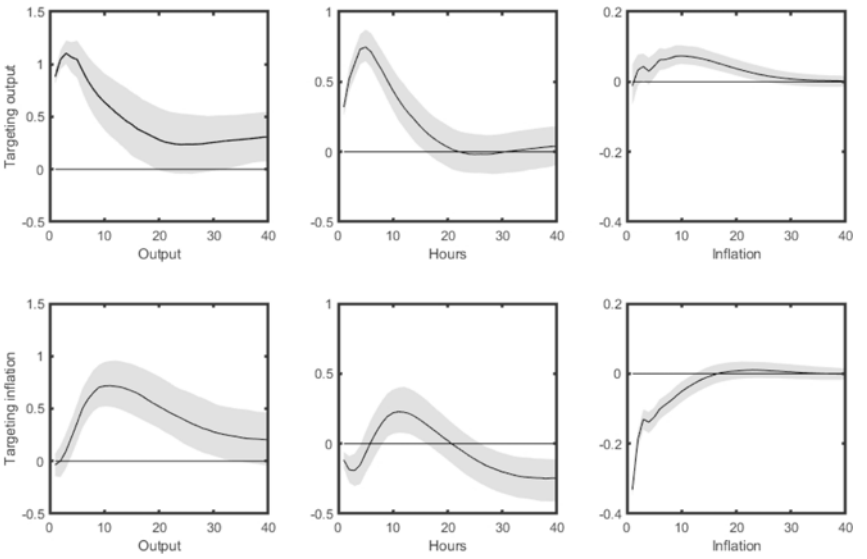



**Figure 2.** Targeting output and inflation at business-cycle frequencies.
*Note:* $16^{th}$ and $86^{th}$ percentile error bands. The first row shows impulse responses to the shock that maximizes the variance of output at business-cycle frequencies (6–32 quarters). The second targets inflation. The VAR consists of US data on labor productivity, hours-worked, the investment to GDP ratio, the consumption to GDP ratio, the consumption expenditure deflator (inflation), and the 10-year US treasury yield.

### *5.3 Identifying the Business-Cycle Anatomy with US Data: the Dichotomy Revealed*

The Spectral approach to Max-Share has been recently applied to finding dominant business-cycle frequency shocks. For example, Angeletos et al. (2020) find a single primary driver of multiple macroeconomic variables at business cycle frequencies.

Using the Spectral method, Angeletos et al. (2020) argue that dominant business-cycle drivers of real variables such as GDP, unemployment, and investment have very little impact on inflation, thereby ruling out inflationary demand shocks of the standard New Keynesian variety as key drivers of the business cycle. We estimate the same VARs used to evaluate the identification of technology in the US data and confirm this finding. Figure 2 presents the impulse responses. By targeting the shock that maximizes the variance of GDP at business cycle frequencies (6–32 quarters), the shock found to boost output by 1% in the short run causes inflation to rise by just 0.05 percentage points at its peak. In contrast, the shock that maximizes the variance of inflation at business cycle frequencies has a more substantial impact on inflation (−0.3 percentage points) but a smaller impact on output. These findings confirm the business cycle dichotomy highlighted in Angeletos et al. (2020).

### *5.4 Dissecting the Business-Cycle Anatomy with a Simple New Keynesian Model*

Can the dichotomous empirical finding be an artifact of the variance-maximizing method in general, rather than a feature of the true DGP which we hope that our spectral restrictions would capture? For a reasonable parameterization of the model, we argue that an alternative candidate explanation for the absence of significant inflationary impact from the main driver of output at business cycle frequencies is the presence of competing demand and supply-side influences. The combination of these competing forces dulls the response of inflation. Demand-side drivers of output will cause inflation to rise, while supply-side drivers of output can cause output to rise while reducing inflation (as with technology and markup shocks, for example). The presence of both types of shock in driving output will result in a hybrid IRF to inflation, consisting of a mix of the effects of both shocks.

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

A simple New Keynesian model demonstrates how the Spectral methodology can produce these results, even in the case where inflationary demand shocks drive the majority of the variance of output:

$$y_t = \frac{-1}{\sigma}\left(R_t - E_t[\pi_{t+1}]\right) + E_t[y_{t+1}] + \eta_t \tag{15}$$

$$\pi_t = \kappa MC_t + \beta E_t[\pi_{t+1}] \tag{16}$$

$$MC_t = (\sigma + \chi)y_t - (1 + \chi)a_t \tag{17}$$

$$R_t = (\phi_t)y_t + \phi_\pi \pi_t, \tag{18}$$

where $y_t$ is output, $R_t$ is the nominal interest rate, $\pi_t$ is inflation, and MC is marginal costs. $\eta_t$ represents a demand (preferences) shock, while $a_t$ reflects a supply-side technology shock. $\sigma$ is the intertemporal elasticity of substitution, $\chi$ is the Frisch elasticity of labor supply. $\kappa$ is the slope of the Phillips curve and is a function of the probability of not being able to reset prices each period ($\theta$) and the discount rate ($\beta$) : $\kappa = (1-\theta)(1-\beta\theta)/\theta$.

In this standard framework, there are two key drivers of fluctuations in output and inflation. Let us suppose that these are transitory drivers, with sizable contributions to the business cycle. These shocks therefore follow auto-regressive processes with a coefficient of 0.7 on the shock's auto-regressive coefficients ($a_t = 0.7a_{t-1} + \epsilon_a$ and $\eta_t = 0.7\eta_{t-1} + \epsilon_\eta$). Elasticities are set to 1 for simplicity ($\sigma$ and $\chi$) while the inflation aversion coefficient is set to 1.5 and the output gap coefficient is set to 0.5 in the policy-rate equation. Shock variances are set at $\sigma_a^2 = 0.004$ and $\sigma_\eta^2 = 0.002$, which results in the supply shock explaining roughly three-quarters of the variance of inflation and the demand shock explaining roughly three-quarters of the variance of output.

Simulating the data for 250 periods and then separately identifying the dominant business-cycle driver of output and inflation in a two-variable VAR using the Spectral approach, as in Angeletos et al. (2020), produces the postulated result, presented in Figure 3. Despite the presence of clearly separate driving processes, the IRFs for inflation end up as a mixture between the supply and demand shock IRFs. This provides the appearance of a subdued response of inflation to a demand shock, despite the demand shock dominating the variance of output in the true DGP. In addition, the response of output to the supply shock is subdued due to the offsetting impact of the demand shock.[4]

The Spectral method will erroneously conclude that a dominant structural demand shock has little effect on inflation. Practitioners will need to have an *a priori* understanding of the range of potential confounding shocks in these exercises and interpret the results accordingly.

**Additional restrictions within the band.** One might ask whether hybridization at business-cycle frequencies could be reduced by imposing additional moment conditions – such as covariance restrictions, sign restrictions, or exclusion restrictions – within the 6–32 quarter band. While such restrictions can in principle help separate demand and supply shocks (see Francis and Kindberg-Hanlon (2022), they go beyond variance maximization and introduce auxiliary identifying structure. Our objective in this section is instead to show that variance-based identification alone, even when targeted to the business-cycle band, cannot recover pure structural shocks when the underlying shocks have overlapping spectral support. This limitation is thus inherent to the informational content of the data, not to the particular estimator we use.

## 6. Empirical conditions under which max-share aligns with ist news

A natural benchmark for evaluating the variance–maximizing identification approach is provided by empirical work showing that, under specific conditions, Max-Share can recover a technologically interpretable shock. In particular, Ben Zeev (2018, 2019) demonstrate that when

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

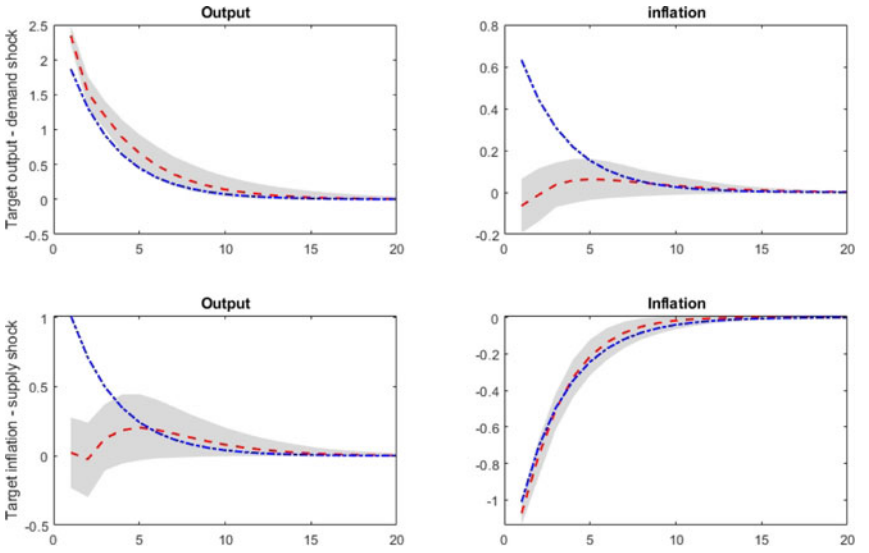



**Figure 3.** Targeting output and inflation at business-cycle frequencies in a two shock New Keynesian model.
*Note:* $16^{th}$ and $86^{th}$ percentile error bands. The first row shows impulse responses to the shock, which maximizes the variance of output produced by the simple New Keynesian model at business-cycle frequencies (6–32 quarters). The second targets inflation. Blue lines indicate the true IRF of the demand shock (top row) and supply shock (bottom row). The shaded area and red line indicate the VAR-estimated IRF median and 68% confidence bands.

identification incorporates the *equipment-based* relative price of investment (RPI) and imposes the agnostic restriction that its long-run behavior is driven by two shocks – an unanticipated IST disturbance and an IST news component—the variance–maximizing shock behaves empirically like IST news. When considered jointly, these findings suggest that the shock accounting for the bulk of business-cycle fluctuations may coincide with an IST news shock under a specific set of empirical conditions.

To assess the robustness of this mechanism in our data, we estimate a structural VAR augmented with an equipment-based relative price of investment series. The data are drawn from the macroeconomic dataset compiled in Ramey (2016) and include quarterly observations on real GDP, total hours, real consumption, real nonresidential investment, and the equipment relative price of investment.[5] The VAR is estimated with four lags and a constant over the sample period 1954:Q1–2018:Q4.

Let $y_t$ denote the $k \times 1$ vector of observables in the VAR,

$$y_t = A_1 y_{t-1} + \cdots + A_p y_{t-p} + u_t, \tag{19}$$

where $u_t$ are reduced-form innovations with covariance matrix $\Sigma_u = \mathbb{E}[u_t u_t']$. Let $C$ be the Cholesky factor such that $\Sigma_u = CC'$, and let $A_0 = CS$ denote the structural impact matrix, where $S$ is an orthonormal matrix that rotates the reduced-form shocks into structural shocks, $\varepsilon_t = S'C^{-1}u_t$.

The moving-average representation is given by

$$y_t = \sum_{h=0}^{\infty} D_h u_{t-h}, \qquad D_0 = I_k, \qquad D_h = A_1 D_{h-1} + \cdots + A_p D_{h-p}. \tag{20}$$

Long-run effects of the structural shocks are summarized by the matrix

$$\Theta(1) \equiv \sum_{h=0}^{\infty} D_h A_0, \tag{21}$$

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

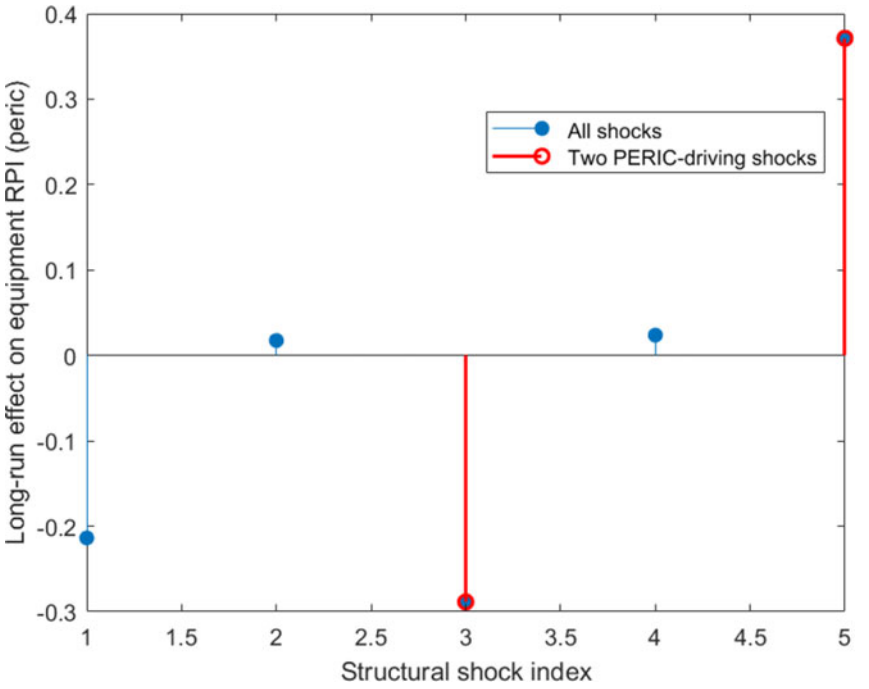



**Figure 4.** Long-run effects of structural shocks on the equipment relative price of investment. Only two structural shocks have non-negligible long-run effects on the equipment RPI, consistent with the imposed agnostic two-shock restriction.

whose $j$th row measures the long-run response of the $j$th variable to each structural shock.

Following Ben Zeev (2018, 2019), we impose an agnostic long-run restriction on the equipment relative price of investment. Let $r$ index the equipment RPI equation. The restriction requires that only two structural shocks – interpretable as an IST surprise and an IST news component – may affect the equipment RPI in the long run:

$$\Theta_{r\ell}(1) = 0 \qquad \text{for all } \ell \notin \mathcal{J}, \qquad |\mathcal{J}| = 2. \tag{22}$$

That is, the $r$th row of $\Theta(1)$ contains exactly two unrestricted elements, while all remaining entries are restricted to zero.

Within the two-dimensional subspace spanned by the shocks in $\mathcal{J}$, we construct the Max-Share shock as the linear combination that maximizes the share of the $K$-step-ahead forecast error variance of output. Let $e_1$ denote the unit vector selecting output, and let $\alpha$ be a $2 \times 1$ vector representing a linear combination of the two RPI-driving structural shocks. The Max-Share problem is

$$\max_{\alpha:\|\alpha\|=1} \quad \frac{\sum_{h=0}^{K} \left(e_1' D_h A_{0,\mathcal{J}} \alpha\right)^2}{\sum_{h=0}^{K} e_1' D_h \Sigma_u D_h' e_1}. \tag{23}$$

The resulting shock,

$$s_t = \alpha' \varepsilon_{\mathcal{J},t}, \tag{24}$$

is the variance–maximizing linear combination of the two structural shocks that drive the long-run behavior of the equipment relative price of investment.

Figure 4 illustrates the long-run effects of all structural shocks on the equipment relative price of investment. The imposed restriction successfully isolates exactly two shocks that load onto the equipment RPI in the long run.

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

**Table 2.** Pairwise correlations among Max-Share shocks and IST news (Equipment RPI)

| | Ben Zeev IST News | PERIC-Anchored MS | Limited-Spectral MS |
|---|---|---|---|
| Plain Max-Share (MS) | 0.391 | 0.463 | 0.996 |
| PERIC-Anchored Max-Share (MS) | 0.254 | — | 0.430 |
| Limited-Spectral Max-Share | 0.387 | 0.430 | — |

*Notes:* The table reports pairwise correlations among (i) the unrestricted variance–maximizing shock (Plain Max-Share), (ii) the Max-Share shock restricted to the two structural shocks that drive the long-run behavior of the equipment-based relative price of investment (PERIC-Anchored Max-Share), (iii) the Limited-Spectral Max-Share shock maximizing low-frequency variance, and (iv) the IST news series constructed by Ben Zeev (2018). Sample period: 1954:Q1–2018:Q4.

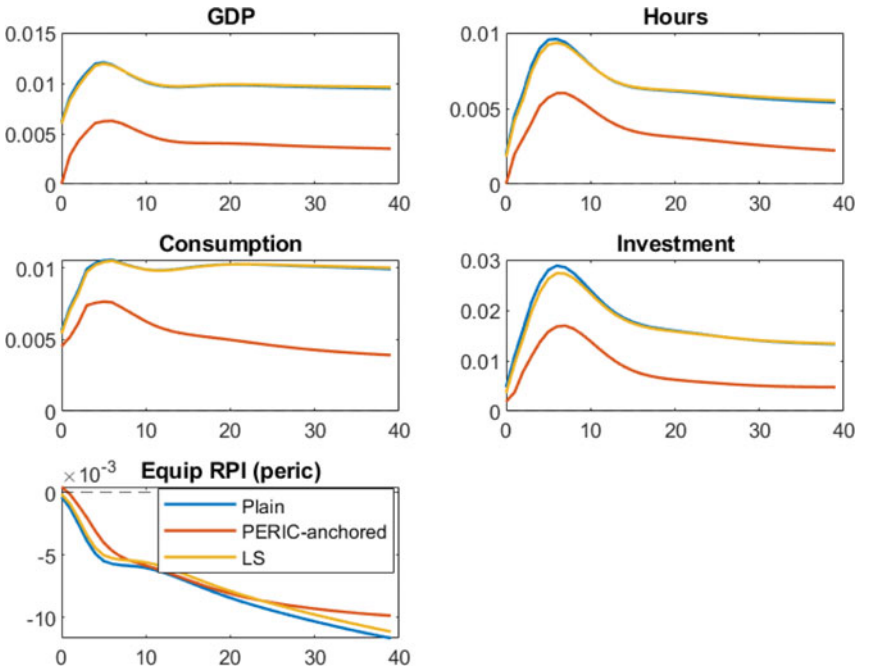



**Figure 5.** Impulse responses of GDP, hours, consumption, investment, and the equipment relative price of investment to three Max-Share shocks: unrestricted (Plain), equipment-price–anchored (PERIC-Anchored), and limited-spectral. The unrestricted and limited-spectral shocks produce nearly identical dynamics, while the PERIC-anchored shock generates distinct responses.

We then compare this restricted Max-Share shock to (i) the unrestricted Max-Share shock obtained from the full $k$-dimensional shock space, (ii) a Limited-Spectral Max-Share shock that maximizes output variance over low-frequency bands, and (iii) the IST news series constructed by Ben Zeev (2018).

Despite imposing this structure, the resulting equipment-price–anchored Max-Share shock does not converge to IST news. The correlation between the anchored Max-Share shock and Ben Zeev's IST news series is 0.254. By contrast, the unrestricted Max-Share shock and the Limited-Spectral Max-Share shock are nearly identical, with a correlation of 0.996, and both exhibit substantially higher correlations with IST news (0.391 and 0.387, respectively). These correlations are summarized in Table 2.

Impulse responses reported in Figure 5 reinforce this conclusion. The unrestricted and Limited-Spectral Max-Share shocks generate nearly identical dynamics for output, hours, consumption, and investment. In contrast, the equipment-price–anchored Max-Share shock produces qualitatively distinct responses, despite being constructed from the shocks that drive the long-run behavior of the equipment price.

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

These findings qualify the interpretation that variance–maximizing identification generically recovers IST news once the equipment-based relative price of investment is used. While Ben Zeev (2018, 2019) demonstrate that such an equivalence can arise under specific empirical conditions, our results show that it is not robust even when the same identifying logic is imposed. Accordingly, the empirical success of Max-Share identification should be interpreted as context-dependent rather than generic.

The near-identity of the unrestricted Max-Share and Limited-Spectral shocks in this application indicates that the forces confounding variance-maximizing identification operate predominantly within the low-frequency band itself, rather than at higher frequencies. In such an environment, frequency restriction does not alter the maximization problem, and both approaches recover the same hybrid combination of low-frequency shocks. By contrast, the PERIC-anchored Max-Share shock differs because the imposed long-run restriction changes the space over which variance maximization is conducted, forcing the estimator to choose a linear combination within the two-shock subspace that drives the long-run behavior of the equipment price. While this restriction successfully isolates shocks with permanent effects on the relative price of investment, it does not eliminate confounding within that subspace: multiple low-frequency forces continue to contribute to output variance, so the resulting variance-maximizing shock need not align with IST news. Therefore, when confounding arises within the targeted frequency range, neither frequency restriction nor long-run anchoring alone is sufficient to recover a pure structural shock.

## 7. Conclusion

This paper documents how confounding shocks can distort variance–maximizing identification methods such as Max-Share, particularly in long-horizon applications. We show analytically why Max-Share recovers hybrid shocks rather than the true dominant long-run driver, and we introduce a frequency-domain alternative – the Limited-Spectral approach – that downweights shocks operating outside the targeted low-frequency band. Through a simple two-shock model and extensive Monte Carlo experiments, we demonstrate that the Limited-Spectral method better isolates the intended low-frequency component.

However, we also highlight an important limitation: when multiple shocks contribute materially within the same frequency range, as is typical at business-cycle frequencies, even spectral variance–maximization can produce hybrid shocks whose interpretation requires additional identifying structure. This underscores a broader point: while frequency-domain refinements can mitigate confounding in environments with clear spectral separation, variance-based identification alone cannot disentangle shocks when their spectral representations overlap.

Finally, we examine whether Max-Share can recover an interpretable IST news shock once it is combined with empirically relevant, yet agnostic, long-run restrictions. Ben Zeev (2018, 2019) show that when identification incorporates the equipment-based relative price of investment and explicitly recognizes its two long-run IST drivers, the variance–maximizing shock can behave empirically like IST news. We implement this identifying logic directly by including an equipment-based relative price of investment series and imposing the same two-shock long-run restriction.

While this restriction is empirically meaningful in our data and isolates the shocks that permanently affect the equipment price, the resulting RPI-anchored Max-Share shock remains only weakly correlated with IST news. At the same time, the unrestricted and Limited-Spectral Max-Share shocks remain nearly identical to one another and continue to differ substantially from IST news. These findings indicate that the empirical equivalence between variance-maximizing identification and IST news documented in earlier work is not generic, even when the equipment-based price and its long-run behavior are explicitly incorporated.

The IST application thus illustrates a broader principle emphasized throughout the paper. Variance-based identification can yield economically interpretable shocks in environments where the data and the imposed restrictions jointly deliver a sufficiently distinctive long-run structure.

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

When this condition is not met, variance maximization does not isolate a single technological disturbance but instead produces hybrid shocks that remain confounded with other structural forces. Accordingly, the empirical success of Max-Share identification should be viewed as context-dependent rather than as a general validation of variance-maximizing approaches.

**Acknowledgements.** We would like to thank seminar participants at Carleton University, the European Central Bank, Penn State, UC Berkeley, UCSD, West Virginia University, the Cleveland Fed, the 2021 poster session at the 45th NBER-NSF Time Series Conference, and the Midwest Macro Conference, Atlanta 2019. We have also benefited from conversations with Michal Andrle, Calebe Figueiredo, James Hamilton, Simona Manu, Garey Ramey, Valerie Ramey, Michele Lenza, and Robert Vigfusson. The views expressed are those of the authors and do not necessarily reflect those of the European Central Bank, the World Bank, its Executive Directors, or the countries they represent.

## Notes

**1** We refer the reader to the Appendix for further details.
**2** In practice we flip the sign of $\hat{s}_t$ and of the associated IRFs whenever the inner product between the estimated and true $L_t$ IRFs is negative.
**3** The differences are more pronounced, having more of the confidence bands non-overlapping, when $T$ is increased.
**4** While our exercise is conducted using the Spectral identification, the finding is identical when we use the Limited-Spectral or Max-Share. For Max-Share, a horizon of 5 years, falling between the typical business-cycle bounds of 2–8 years, displays the same IRF bias and has the same correlation with the true underlying shocks as the Spectral method.
**5** The equipment price series, Relative Price of Equipment, Index 2017 = 1, Quarterly, Seasonally Adjusted (peric), is obtained from the FRED database.

## References

Angeletos G.-M., F. Collard and H. Dellas. (2020). Business-cycle anatomy. *American Economic Review* 110(10), 3030–3070.
Ben Zeev N. (2018). What can we learn about news shocks from the late, 1990s, and early, 2000s, boom–bust period? *Journal of Economic Dynamics and Control* 87, 94–105.
Ben Zeev N. (2019). *Is there a Single Shock that Drives the Majority of Business Cycle Fluctuations?* Ben-Gurion University of the Negev, Department of Economics. Working Paper 1906.
Chari V., P.J. Kehoe and E.R. McGrattan. (2008). Are structural VARs with long-run restrictions useful in developing business cycle theory? *Journal of Monetary Economics* 55(8), 1337–1352.
Christiano L., M. Eichenbaum and R. Vigfusson. (2003). *What Happens After a Technology Shock?* Cambridge, MA: National Bureau of Economic Research. NBER working paper.
Dieppe A., N. Francis and G. Kindberg-Hanlon. (2021). *Technology and Demand Drivers of Productivity Dynamics in Developed and Emerging Market Economies*. Frankfurt: European Central Bank. ECB Working Paper 2533.
Dou L., P. Ho and T. Lubik. (2025). Max-Share Misidentification. *Federal Reserve Bank of Richmond Working Papers* 25(02), 1–36. Working Paper Series 25-2, Federal Reserve Bank of Richmond. https://doi.org/10.21144/wp25-02
Francis N. and G. Kindberg-Hanlon. (2022). Signing out confounding shocks in variance-maximizing identification methods. *AEA Papers and Proceedings* 112, 476–480.
Francis N., M.T. Owyang, J.E. Roush and R. DiCecio. (2014). A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks. *Review of Economic Statistics* 96(4), 638–647.
Galí J. and P. Rabanal. (2005). Technology shocks and aggregate fluctuations: How well does the real business cycle model fit postwar U.S. data? In *NBER Macroeconomics Annual 2004*. Cambridge, MA: National Bureau of Economic Research, pp. 225–318.
Giannone D., M. Lenza and L. Reichlin. (2019). Money, credit, monetary policy, and the business cycle in the Euro area: What has changed since the crisis? *International Journal of Central Banking* 15(5), 137–173.
Ramey V.A. (2016). Chapter 2 - Macroeconomic shocks and their propagation. In Taylor, J.B. and Uhlig, H. (eds.), *Handbook of Macroeconomics*. Elsevier, Vol. 2, pp. 71–162.
Uhlig H. (2003). *What Moves Real GNP?* mimeo HU Berlin

## Appendix A. Solving the maximization problem

Even where the low-frequency shock dominates the forecast error at horizon $k$, the shock captured by Max-Share will not consist only of this underlying shock. Here, we show that the shock that

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

maximizes the contribution to the forecast error variance of productivity is a combination of the low and high-frequency structural shock rather than the dominant underlying structural shock.

Start with the definition

$$Y_t = \begin{bmatrix} p_t \\ n_t \end{bmatrix} = B(L)u_t.$$

Productivity ($p$) and hours ($n$) are a function of the reduced form MA-coefficients (B(L)) and errors ($u$). The true underlying structural shocks are $A_0^{-1}u = \epsilon = \begin{bmatrix} \epsilon_l \\ \epsilon_h \end{bmatrix}$, where $\epsilon$ refers to a low and high-frequency shock. If an empiricist wanted to identify the low-frequency driver of productivity, they may employ the Max-Share approach, assuming that the low-frequency driver will provide the largest contribution to the forecast error variance over a sufficiently long time horizon. In the Max-Share approach, the search for $A_0$ begins with $\tilde{A}_0$, a Cholesky decomposition of $\Sigma_u$ and the impulse matrix $\tilde{A}^\tau = D^\tau \tilde{A}_0$ for the impact of the shock at horizon $\tau$. The "true" structural shock is defined using an unknown orthonormal matrix $Q$, such that the true structural impulse matrix $A^\tau = \tilde{A}^\tau Q$ and the forecast error at horizon $k$ is

$$F(Y)_k = \sum_{\tau=0}^{k-1} \tilde{A}^\tau Q \epsilon_{t+k-l} = \sum_{\tau=0}^{k-1} A^\tau \epsilon_{t+k-\tau}$$

with the error variance also equivalent for both the true shock and the initially proposed Cholesky decomposition form since:

$$\Sigma(k) = \sum_{\tau=0}^{k-1} \tilde{A}^\tau \tilde{A}^{\tau\prime} = \sum_{\tau=0}^{2} \tilde{A}^\tau Q Q' \tilde{A}^{\tau\prime}$$

As in Uhlig (2003), isolating the shock that explains the largest proportion of the forecast error variance reduces to an eigenvector decomposition problem. This means searching for an orthonormal vector $q$ (of the matrix $Q$) which maximizes the forecast error variance for variable *ii*.

$$\begin{aligned} F^2(Y)_k &= \sum_{\tau=0}^{k-1} \left( (D^\tau \tilde{A}_0 q)(D^\tau \tilde{A}_0 q)' \right)_{ii} \\ &= q' S(k) q, \end{aligned}$$

where

$$S(k) = \sum_{l=0}^{k-1} (D_i^\tau \tilde{A}_0)(D_i^\tau \tilde{A}_0)'.$$

Maximizing $q'S(k)q$ is a Lagrangian with the constraint that $q'q = 1$

$$L(q) = q' S(k) q - \lambda (q'q - 1)$$

whose first order conditions reduce to solving for the eigenvector associated with the largest eigenvalue of $S(k)$

$$S(k)q = \lambda q.$$

Will this $q$ be equivalent to the first column of $Q$ defining the true low and high-frequency structural shocks, $A^0 = \tilde{A}^0 Q$?

We will demonstrate that, in an attempt to isolate the low-frequency shock, the $q$ in the Max-Share rotation will still contain features of the high-frequency shock, the extent depending on the

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

forecast horizon, the relative variances of the shocks, and the relative persistence of the shocks. While we demonstrate this for the simple 2 × 2 case, we argue that the results hold for higher dimension VARs.

The true underlying low and high-frequency shocks can be written as

$$A_0\epsilon = \begin{bmatrix} A_{lp} & A_{hp} \\ A_{ln} & A_{hn} \end{bmatrix} \begin{bmatrix} \epsilon^l \\ \epsilon^h \end{bmatrix}$$

Here, $A_{lp}$ and $A_{ln}$ refer to the impacts of the low-frequency shock on productivity and hours respectively in period 0, while $A_{hp}$ and $A_{hn}$ refer to the impacts of the high-frequency shock on productivity and hours respectively.

The total forecast error variance is equivalent when using the structural form of the shocks and the reduced form since

$$\Sigma_u = A_0 A_0'$$

The forecast error variance can be computed using the structural shocks and the reduced form MA-representation coefficients $D_{2\times 2}^{i:T}$. The empiricist will initially use the Cholesky decomposition of $\Sigma_u$ as a starting point for the identification of the true shocks. Assuming unit variances and orthogonality of the structural shocks ($E(\epsilon^2) = 1$ and $E(\epsilon^l \epsilon^h) = 0$):

$$Chol(A_0 A_0') = \tilde{A}_0 = \begin{bmatrix} \sqrt{A_{lp}^2 + A_{hp}^2} & 0 \\ \frac{A_{lp}A_{ln} + A_{hp}A_{hn}}{\sqrt{A_{lp}^2 + A_{hp}^2}} & \frac{A_{ln}A_{hp} - A_{hn}A_{lp}}{\sqrt{A_{lp}^2 + A_{hp}^2}} \end{bmatrix}$$

In the maximization problem, we are only interested in the impacts on the first endogenous variable of the system (productivity).

$$F^2(Y)_K = \sum_{i\tau=0}^{k-1} \left( \left[ D_\tau^{11} D_\tau^{12} \right] \tilde{A}_0 \right)' \left( \left[ D_\tau^{11} D_\tau^{12} \right] \tilde{A}_0 \right)$$

at $K = 0$, $[D_0^{11}\ D_0^{12}] = [1\ 0]$ since the second variable will not have a contemporaneous impact on the first.

Therefore, in this very simple case:

$$F^2(Y)_0 = \begin{bmatrix} A_{lp}^2 + A_{hp}^2 & 0 \\ 0 & 0 \end{bmatrix}$$

As in Uhlig, finding the shock that maximizes the forecast error variance of variable $i$ reduces to an eigenvector problem. In the case of $S(0) = FEV(0)$, the eigenvalue problem is simple. Here, there are two eigenvalues found, $\lambda = [A_{lp}^2 + A_{hp}^2, 0]$. The eigenvector associated with the largest eigenvalue, the non-zero value, is simply:

$$\tilde{q}(0) = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

Therefore, in this simple case of maximizing the forecast error variance at period zero, the dominant shock is solved as:

$$\tilde{A}\tilde{q}(0) = \begin{bmatrix} \sqrt{A_{lp}^2 + A_{hp}^2} \\ \frac{A_{lp}A_{ln} + A_{hp}A_{hn}}{\sqrt{A_{lp}^2 + A_{hp}^2}} \end{bmatrix}$$

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press

The two entries are the impact of the shock on productivity and hours respectively. Note that the impact on productivity will reflect both the low and high-frequency "true" shocks in proportion to their scale. Note also that the impact on hours will reflect both shocks.

In most cases of applying the Max-Share methodology, the practitioner will maximize contributions to the forecast error variance over much longer horizons. This will result in more complex forms for $[D^{11}\ D^{12}]$ and a more complex eigenvalue-eigenvector problem. However, the maximization will still capture a combination of underlying structural shocks rather than a single dominant structural shock.

Over longer horizons, the maximization becomes more analytically complex. At horizon $\tau$, we have:

$$h_\tau = [D_\tau^{11} D_\tau^{12}]\tilde{A}_0$$

And

$$h_{\tau,1} = D_\tau^{11}\sqrt{A_{lp}^2 + A_{hp}^2} + D_\tau^{12}\left(\frac{A_{lp}A_{ln} + A_{hp}A_{hn}}{\sqrt{A_{lp}^2 + A_{hp}^2}}\right)$$

$$h_{\tau,2} = D_\tau^{12}\left(\frac{A_{ln}A_{hp} - A_{hn}A_{lp}}{\sqrt{A_{lp}^2 + A_{hp}^2}}\right)$$

with

$$F^2(Y)_K = \sum_{\tau=1}^{k-1}\begin{bmatrix} h_{\tau,1}^2 & h_{\tau,1}h_{\tau,2} \\ h_{\tau,2}h_{\tau,1} & h_{\tau,2}^2 = \end{bmatrix} = \begin{bmatrix} \mathcal{M}_{11} & \mathcal{M}_{12} \\ \mathcal{M}_{21} & \mathcal{M}_{22} \end{bmatrix}$$

The eigenvalues of $F^2(Y)_K$ are now more complex than in the standard case, since the addition of the $h$ terms at each horizon mean that $\mathcal{M}_{12}$ is not necessarily equal to $\mathcal{M}_{11}\mathcal{M}_{22}$. If this *was* still the case then the maximization problem would be similar to the period 0 case, and the eigenvalues would be $[\mathcal{M}_{11} + \mathcal{M}_{22}, 0]$. The eigenvector associated with the largest eigenvalue would be $\begin{bmatrix} \frac{\mathcal{M}_{11}}{\mathcal{M}_{22}} \\ 1 \end{bmatrix}$. Therefore, the shock identification in this case is a complex function of the size of the initial impacts of each of the true shocks ($A_0$) and the persistence of their impacts over the chosen horizon.

In practice, the matrix $F^2(Y)_K$ would have full rank and a more complex structure, but would nonetheless have the same properties in combining a component of each shock that was a function of the impact variance and IRF frequency properties.

**Cite this article:** Dieppe A, Francis N and Kindberg-Hanlon G (2026). "Variance-Maximizing Identification and the Plague of Confounding Shocks." *Macroeconomic Dynamics* **30**(e28), 1–22. https://doi.org/10.1017/S1365100526101023

https://doi.org/10.1017/S1365100526101023 Published online by Cambridge University Press