Max-Share Misidentification*

Liyu Dou[1], Paul Ho[2], and Thomas A. Lubik[3]

[1]Singapore Management University
[2,3]Federal Reserve Bank of Richmond

February 11, 2026

Abstract

While max-share identification has become increasingly popular in a wide range of applications, we show that its validity requires necessary and sufficient conditions that are rarely satisfied in practice—the target variable's response to the target shock must be (i) orthogonal to its responses to untargeted shocks and (ii) larger than combinations of those responses. Imposing additional restrictions on the target shock weakens but does not fully eliminate these conditions. We show that in practice, the weight max-share places on an identified untargeted shock can be obtained by projecting the response to that shock on the max-share response. We also theoretically characterize consequences of local and global violations to the identification conditions. Empirically, the TFP news and main business cycle shocks identified by Kurmann and Sims (2021) and Angeletos et al. (2020) are, respectively, at least a third and a quarter contaminated.

*We are grateful to Drew Creal, Lutz Kilian, Michal Kolesár, Peter Phillips, Mikkel Plagborg-Møller, Felipe Schwartzman, and seminar and conference participants at the Barcelona Summer Forum, Conference of the International Association for Applied Econometrics, International Symposium on Econometric Theory and Applications, Econometric Society World Congress, RCEA International Conference, and University of Illinois Urbana-Champagn for comments and suggestions. Katie Anderson provided excellent research assistance. The views expressed herein are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Richmond or the Federal Reserve System.
Liyu Dou: liyudou@smu.edu.sg.
Paul Ho: paul.ho@rich.frb.org.
Thomas A. Lubik: thomas.lubik@rich.frb.org.

1 Introduction

Max-share identification has become a popular approach for identifying structural shocks in vector autoregressions (VARs), most prominently a total factor productivity (TFP) news shock (Beaudry and Portier, 2006; Barsky and Sims, 2011; Kurmann and Sims, 2021) and a main business cycle shock (Angeletos, Collard, and Dellas, 2020). Max-share identifies a structural shock as the one that maximizes its contribution to a particular economic variable's forecast error variance (FEV) over some horizon (in the time domain) or variation at some frequency (in the frequency domain). This approach is attractive because it yields point identification while appearing to rest only on the innocuous assumption that the target shock has a larger effect than other shocks on the target variable at the chosen horizon or frequency.

We show that max-share identification relies on stringent conditions to be valid. We derive necessary and sufficient conditions and demonstrate that they are considerably harder to satisfy than previously recognized in the literature. These conditions rely on the fact that max-share identification is equivalent to obtaining the principal eigenvector of a matrix, which we denote by Ξ. This matrix can be interpreted heuristically as capturing the covariance across impulse responses of the target variable to the true structural shocks. Instead of being of random variables over observations, the covariance is of impulse responses over horizons.

The matrix, Ξ, is a function of the reduced-form parameters, which we assume to be known, and the frequency or horizon of interest. Its form depends on whether we implement max-share identification in the time or frequency domain. In the time domain, Ξ captures the FEV decomposition for the true shocks; in the frequency domain, Ξ captures the contribution of each shock to the variation of the target variable over the chosen frequency band. The diagonal terms capture the magnitude of the impulse responses over the target set of horizons or frequency band; the off-diagonal terms capture the similarity in the shape or lead-lag structure over those horizons or frequencies. Since our identification conditions are restrictions on Ξ, which, in turn, is directly connected with the size and shape of the impulse responses, the identified shock will be contaminated unless strict conditions on the impulse responses to the full set of true shocks are also satisfied. This contrasts with the typical justification given in the literature about the relative magnitude of a response at a particular horizon or frequency. In what follows, we assume without loss of

generality that the true target shock is ordered first for ease of exposition.

The first identification condition is *orthogonality* of the impulse responses. This corresponds to Ξ being a block diagonal matrix with the off-diagonal terms in the first row and column being zero. It is a restriction on the shape of the target impulse response relative to the response to each of the other shocks over the target horizons or frequencies. In the time domain, the stringency of the orthogonality condition is apparent. Suppose we use max-share for some horizon H. Then, noting that the FEV sums the contributions of shocks from horizons 0 to H, we can represent the impulse response to the jth true shock as a $(H+1) \times 1$ vector, ψ_j, where the hth element of ψ_j corresponds to the response of the target variable to shock j at horizon $h = \mathsf{h} - 1$. Orthogonality requires that $\psi_1 \cdot \psi_j = 0$ for all $j \neq 1$. This rules out, for instance, cases in which the target shock and some untargeted shock both produce strictly positive impulse responses. In the frequency domain, orthogonality is similarly violated in many standard models. Since orthogonality has to be satisfied for the full set of shocks, it requires a priori knowledge of how the target variable responds to all the shocks over a given set of horizons, which is arguably more demanding than what is required for zero or sign restrictions.

The second condition is on the *relative size* of the target impulse response. Formally, the $(1,1)$ element of Ξ must be larger in magnitude than the largest eigenvalue of the lower $(N-1) \times (N-1)$ block of Ξ, where N is the total number of shocks in the VAR. Intuitively, this means that the target impulse response must be large relative to not only the response to each of the other individual shocks, but also combinations of these other impulse responses. For instance, suppose the responses to each of the untargeted shocks are small. If these responses have similar shapes, then the eigenvalue of the lower $(N-1) \times (N-1)$ block of Ξ can still be large and cause the max-share identified shock to be a combination of the untargeted shocks. We also note that the magnitude of each of the elements of Ξ will depend on the size of the impulse response over a set of horizons. In the time domain, the FEV at horizon H depends on the cumulative contribution of a shock over horizons 0 to H. Hence, the magnitude of the corresponding impulse response and diagonal element in Ξ depends on the entire response over periods $0, \ldots, H$ and not just the size of the response at horizon H. In the frequency domain, the magnitude depends on the impulse response over all horizons $h \geq 0$.

As a tool for empirical research, we propose a straightforward diagnostic that

yields a lower bound on how much the max-share identified shock is contaminated in practice. Suppose a researcher has information on the impulse response of the target variable to an untargeted shock, for instance, through prior knowledge or by separately identifying that shock. The weight that max-share places on the untargeted shock is the projection of the impulse response of the target variable to the untargeted shock on the corresponding response to the max-share shock. Consequently, we can place an upper bound on the max-share weight on the true targeted shock. In the presence of multiple untargeted shocks this upper bound is generally conservative so that the actual contamination of the max-share identified shock is even worse.

The identification conditions can also be adapted to versions of max-share with additional linear constraints. This nests approaches such as zero restrictions in the news shock literature (e.g., Barsky and Sims, 2011) or controlling for other shocks (e.g., Basu et al., 2025). Instead of the space spanned by the impulse responses of the target variable to untargeted shocks, the orthogonality and relative size conditions now pertain to the subspace that is orthogonal to the constraints. The resulting identification restrictions are weaker but can still be difficult to satisfy.

Finally, we present two theoretical results on the behavior of max-share identification when the identification conditions are violated. First, we consider a perturbation from the identified case that generates small violations of orthogonality. The resulting weight on the untargeted shocks can be expressed, to first order, as the eigenvectors scaled by the projection of the perturbation on the eigenvector divided by the difference between the corresponding eigenvalue and the principal eigenvalue. Second, we show that for an arbitrary deviation in the orthogonality condition, we can bound the weights on untargeted shocks by a function that depends inversely on the difference between the two largest eigenvalues. These results emphasize that when the identification conditions are not strictly satisfied, the performance of max-share depends critically on the true response of the target variable to the target shock being substantially larger than the corresponding responses to all the untargeted shocks.

In light of our results, we revisit the shocks identified by Kurmann and Sims (2021) and Angeletos et al. (2020). Based on the VAR from Kurmann and Sims (2021), we identify a TFP news shock and a main business cycle shock using max-share in the time and frequency domains, respectively. In addition, we identify a TFP surprise shock using recursive identification. These three shocks are typically interpreted as being distinct, with the TFP shocks differing in their initial impact and the main

business cycle shock taken as a demand shock independent of TFP. However, using our diagnostic described above, we find that over a third of the identified TFP news shock comes from the TFP surprise shock and over a quarter of the business cycle shock comes from the TFP news shock. We further show numerical examples of how these identification issues are present in ARMA, demand and supply, and medium-scale New Keynesian (Smets and Wouters, 2007) models.

Related Literature. The max-share problem was initially introduced in the time domain by Faust (1998) to obtain bounds on VAR impulse response functions. The idea was refined and implemented for identification purposes by Uhlig (2004a) and Uhlig (2004b). Subsequently, DiCecio and Owyang (2010) adapted max-share identification to the frequency domain.

Since these contributions, the max-share approach has been used to identify a vast array of shocks.[1] Our results also nest related work by Barsky and Sims (2011), Kurmann and Otrok (2013), and Ben Zeev and Khan (2015) that maximizes the contribution to the sum of FEVs over various horizons possibly subject to additional constraints. For many of these shocks, it is a challenge to impose appropriate zero and sign restrictions (Sims, 1980; Uhlig, 2005; Arias, Rubio-Ramírez, and Waggoner, 2018) or to find suitable instruments for identification (Mertens and Ravn, 2013; Stock and Watson, 2018). In the literature, the stated underlying assumption is typically that identification only requires that the target shock is important for a particular variable at some horizon or frequency. However, we show that the identification scheme is less innocuous than it initially appears.

A spate of recent research has pointed out potential issues with max-share identification. However, these contributions tend to focus on specific applications rather than presenting a more general case for the deficiency of max-share. For instance, Dieppe, Francis, and Kindberg-Hanlon (2021) raise concerns about the possibility of confounding several shocks when using max-share to identify technology shocks,

[1]These include permanent supply shocks (Francis, Owyang, Roush, and DiCecio, 2014), uncertainty shocks (Caldara, Fuentes-Albero, Gilchrist, and Zakrajšek, 2016), credit shocks (Mumtaz, Pinter, and Theodoridis, 2018), business cycle shocks (Giannone, Lenza, and Reichlin, 2019; Angeletos, Collard, and Dellas, 2020), sentiment shocks (Fève and Guay, 2019; Levchenko and Pandalai-Nayar, 2020), TFP news shocks (Kurmann and Sims, 2021; Görtz, Tsoukalas, and Zanetti, 2022; Görtz, Gunn, and Lubik, 2022), treasury yield news shocks (Moench and Soofi-Siavash, 2022), risk premium shocks (Basu, Candian, Chahrour, and Valchev, 2025), exchange rate shocks (Miyamoto, Nguyen, and Oh, 2023; Chahrour, Cormun, De Leo, Guerrón-Quintana, and Valchev, 2024), and government spending shocks (Chen and Liu, 2019).

which Francis and Kindberg-Hanlon (2022) use sign restrictions to address. Similarly, Kilian, Plante, and Richter (2023) point out that using max-share and related approaches targeting medium horizons to identify news shocks (Barsky and Sims, 2011; Kurmann and Sims, 2021; Dieppe et al., 2021) may produce misleading results and propose using max-share to target direct measures of technological news at short horizons instead. Cascaldi-Garcia and Galvao (2021) show that when used to identify news and uncertainty shocks separately, max-share yields shocks that are highly correlated instead of independent, an issue Carriero and Volpicella (2024) resolve by jointly identifying multiple shocks by max-share. Our general conditions for max-share identification serve as a single lens through which to understand the threats to identification in each application.

A more theoretical critique is presented by Guay, Pelgrin, and Surprenant (2024), who extend arguments from Phillips (1998) to show that the inconsistency of the reduced form impulse responses in the presence of unit or near-unit roots threatens the validity of max-share. Our analysis instead assumes that the reduced form parameters are known. In other words, even without the issues emphasized by Guay et al. (2024), max-share still requires stringent conditions to accurately identify the target shock.

Outline. The rest of the paper is organized as follows. Section 2 illustrates the identification conditions using four stylized examples. Section 3 sets up a common framework for max-share identification in the time and frequency domains. Section 4 presents the necessary and sufficient conditions for valid identification. Section 5 presents tools to diagnose or reduce identification issues and theoretical results on contamination when the identification conditions are violated. Sections 6 and 7 illustrate our insights using numerical examples and an empirical application, respectively. Section 8 concludes. All proofs are in the Online Appendix.

2 Building Intuition from Illustrative Examples

We first present four stylized examples in Figure 1 that demonstrate how max-share identification can fail. In each case, the econometrician seeks to identify Shock 1 and knows that it produces the largest response at horizon 1. A common practice in the literature is to take this knowledge as justification for using max-share with target horizon 1 to identify the shock. Panel A presents the ideal case—max-share perfectly

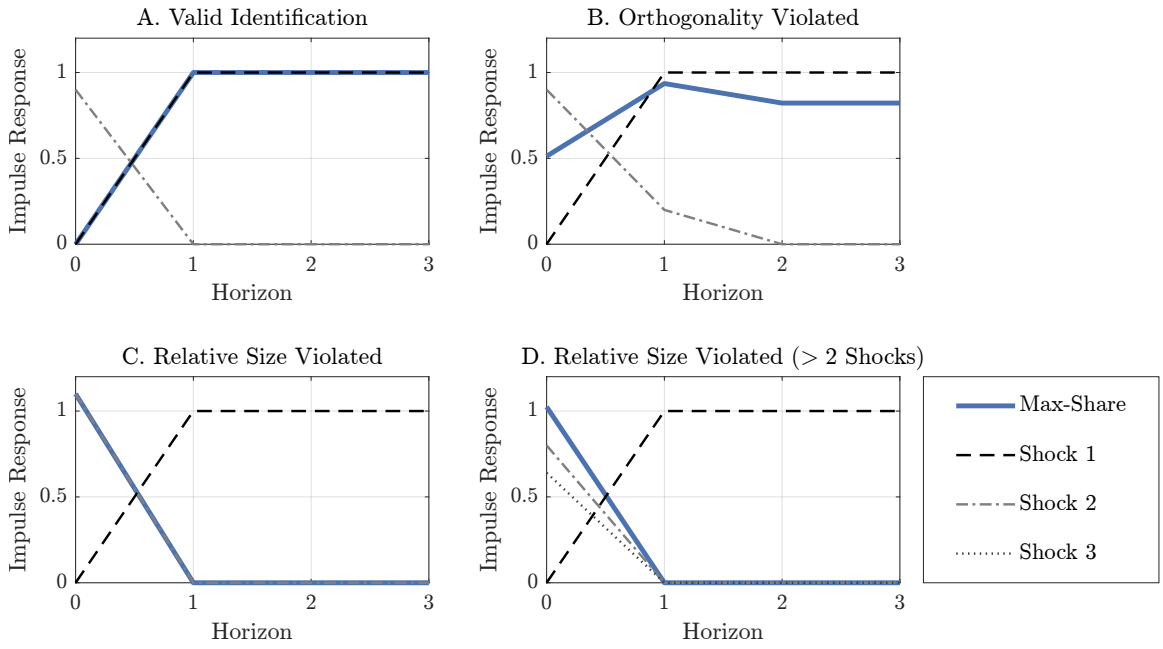


Figure 1: Impulse responses to max-share identified and true shocks in four stylized examples. Solid blue line corresponds to max-share shock; remaining lines correspond to true shocks.

identifies Shock 1. Orthogonality holds because over the horizons 0 and 1, the one other shock, Shock 2, is zero whenever Shock 1 is non-zero and vice versa. Moreover, since the response to Shock 1 peaks at 1 while that to Shock 2 only peaks at 0.9, the relative size condition is also satisfied.

The subsequent panels show how slight deviations threaten identification. In Panel B, orthogonality is violated as we now increase the horizon 1 response to Shock 2 from 0 to 0.2. Even though the response to Shock 1 is still larger ($1 > 0.9^2 + 0.2^2$), the max-share shock is now a combination of Shocks 1 and 2, with less than 60% of the weight on Shock 1. Panel C shows relative size being violated. Returning the horizon 1 response to Shock 2 to 0 so that orthogonality is satisfied, we increase the response on impact to 1.1, making it explain a larger share of the FEV. As a result, the max-share identified shock is incorrectly identical to Shock 2. This emphasizes that earlier horizons have a potentially important role for the relative size condition. Finally, Panel D shows that with more than one untargeted shock, we can violate the relative size condition even if the individual response to each of the untargeted shocks is relatively small. In particular, we reduce the impact response to Shock 2 to 0.8 but introduce an additional shock, Shock 3, that produces a response that is

identical up to a scaling of 4/5. The corresponding eigenvalue accounts for the fact that multiple impulse responses are identical up to scale and consequently dominate the one corresponding to Shock 1. Max-share identification thus picks up Shocks 2 and 3 only. We will show that these issues are pervasive.

3 The Max-Share Identification Problem

Consider a general structural VAR:

$$Y_t = \sum_{\ell=1}^{L} B_\ell Y_{t-\ell} + C\varepsilon_t, \tag{1}$$

where Y_t is an $N \times 1$ vector and ε_t is iid over time, with $E[\varepsilon_t] = 0$ and $E[\varepsilon_t \varepsilon_t'] = I$. Under suitable stationarity conditions, we have the moving average representation:

$$Y_t = \sum_{h=0}^{\infty} \Psi_h \varepsilon_{t-h}, \tag{2}$$

where the $N \times N$ matrix Ψ_h summarizes the impulse responses at horizon h. Each column of Ψ_h corresponds to a shock and each row corresponds to an endogenous variable. The estimates of the reduced form VAR provides $\Sigma = CC'$, but not C. Accordingly, we will assume that $\Psi_h \Psi_h'$ is known, but additional restrictions are required to identify Ψ_h. Max-share identification is one way of obtaining the restrictions for one column of Ψ_h, i.e., identifying one of the structural shocks.

3.1 Time Domain

Max-share identification is most commonly performed in the time domain. To identify Ψ_h, it considers the problem:

$$\arg\max_{\theta} \frac{\delta_i' \left[\sum_{h\in\mathcal{H}} \widetilde{\Psi}_h \theta\theta' \widetilde{\Psi}_h'\right] \delta_i}{\delta_i' \left[\sum_{h\in\mathcal{H}} \widetilde{\Psi}_h \widetilde{\Psi}_h'\right] \delta_i} \text{ subject to } \theta'\theta = 1, \tag{3}$$

where θ is the rotation or vector of weights we are solving for, δ_i is a vector with 1 in the ith entry and 0 everywhere else, $\mathcal{H}$ is a set of horizons chosen by the econometrician, and $\widetilde{\Psi}_h$ is an arbitrary rotation of the structural shocks satisfying $\widetilde{\Psi}_h \widetilde{\Psi}_h' = \Psi_h \Psi_h'$.

The typical application of (3) sets $\mathcal{H} = \{0, 1, \ldots, H\}$. This maximizes the contribution of the target shock to the FEV of variable i. The $N \times N$ matrix $\widetilde{\Psi}_h \theta\theta' \widetilde{\Psi}_h'$ in the numerator captures the contribution of the target shock to the overall variance of the full vector of endogenous variables, Y_t, and multiplication by δ_i extracts the contribution to variable i. This is standard practice in the identification of technology shocks (Francis et al., 2014) and news shocks (Barsky and Sims, 2011). Taking $H \to \infty$ corresponds to long-run identification (Blanchard and Quah, 1989). Setting $\mathcal{H} = \{0\}$ and $i = 1$ corresponds to taking C as the lower triangular matrix from the Cholesky decomposition of Σ. If the first variable is an instrument for the shock of interest, then this is internal instrument identification (Noh, 2018; Plagborg-Møller and Wolf, 2021).

We can recast the time domain max-share problem as an eigenproblem, as first discussed in Faust (1998), which helps unify our exposition of the frequency domain problem (the formal proof can be found in the Online Appendix).

Lemma 1. *Solving* (3) *is equivalent to solving:*

$$\arg\max_{\theta} \theta' \left[\sum_{h \in \mathcal{H}} \widetilde{\Psi}_h' \delta_i \delta_i' \widetilde{\Psi}_h \right] \theta \quad \textit{subject to} \quad \theta'\theta = 1. \tag{4}$$

The solution is the principal eigenvector of:

$$\sum_{h \in \mathcal{H}} \widetilde{\Psi}_h' \delta_i \delta_i' \widetilde{\Psi}_h. \tag{5}$$

subject to $\theta'\theta = 1$.

3.2 Frequency Domain

An alternative approach proposed by DiCecio and Owyang (2010) solves the max-share problem in the frequency domain. The motivation for such an approach is that one may only be interested in fluctuations of particular frequencies. For instance, Stock and Watson (1999) argue that analysis of economic fluctuations should focus on business cycle frequencies to avoid contamination by overly high or low frequency fluctuations in the data.

The frequency domain max-share approach solves the following:

$$\arg\max_{\theta} \theta' \text{Re}\left[\int_{\omega\in\Omega} \Upsilon(\omega)\, d\omega\right] \theta \text{ subject to } \theta'\theta = 1, \tag{6}$$

where, denoting the conjugate transpose of a matrix X by $\overline{\mathsf{X}}$,

$$\Upsilon(\omega) \equiv \overline{\Gamma(\omega)}\delta_i\delta_i'\Gamma(\omega), \tag{7}$$

$$\Gamma(\omega) \equiv \sum_{h=0}^{\infty} \widetilde{\Psi}_h e^{-i\omega h}. \tag{8}$$

$\Omega \subseteq [-\pi, \pi]$ is the set of frequencies of interest (typically some band $[\omega_L, \omega_H]$), δ_i is the vector with 1 in its ith position and 0 everywhere else as before, and Γ is the lag polynomial in (2) with the lag operator replaced by $e^{-i\omega}$, often referred to as the *transfer function*.[2] Similar to the time domain, the solution to (6) is the eigenvector associated with the largest eigenvalue of $\text{Re}\left[\int_{\omega\in\Omega} \Upsilon(\omega)\, d\omega\right]$ as shown in Lemma 1.[3]

3.3 Unifying the Time and Frequency Domain Problems

We analyze max-share in a single framework by using the fact that the time and frequency domain problems share a common structure. For exposition, we henceforth take $\widetilde{\Psi}_h = \Psi_h$, so that the jth entry of the solution, θ, to (3) or (6) is the weight that the max-share shock places on the jth true shock.[4] θ_j is also the correlation between the max-share shock and the jth true shock in population.[5]

In both time and frequency domains, we can write the max-share problem as:

$$\arg\max_{\theta} \theta'\Xi\theta \text{ subject to } \theta'\theta = 1, \tag{9}$$

[2] See Priestley (1981) for a classic reference.

[3] For the problem (6) to be well-defined, we require regularity conditions on the impulse responses, Ψ_h. In what follows, we maintain the high-level assumption that $\sum_{h=0}^{\infty} \|\Psi_h\| < \infty$, where $\|\cdot\|$ denotes the Frobenius norm. Under this assumption, Y_t in (2) is weakly stationary and has finite variance. Moreover, $\Gamma_{ij}(\omega)$ is continuous and bounded on Ω by standard arguments. This is a direct application of the Weierstrass M-test and the uniform limit theorem.

[4] This choice of $\widetilde{\Psi}_h$ affects the solution for θ in (3) but does not change the implied impulse responses to the max-share shock. In particular, if we replace $\widetilde{\Psi}_h$ with a $\widetilde{\Psi}_h R$, where R is an arbitrary rotation matrix, the solution θ will be replaced by $R^{-1}\theta$, leaving the max-share shock unchanged. In practice, since the true responses are unknown, a convenient choice is to take $\widetilde{\Psi}_h$ to be the lower triangular matrix from the Cholesky decomposition of Σ.

[5] This follows since $\text{corr}(\sum_{k=1}^{N} \theta_k \varepsilon_{k,t}, \varepsilon_{j,t}) = \sum_{k=1}^{N} \theta_k \text{corr}(\varepsilon_{k,t}, \varepsilon_{j,t}) = \theta_j \text{corr}(\varepsilon_{j,t}, \varepsilon_{j,t}) = \theta_j$.

where Ξ is a *Gram matrix*, i.e.,

$$\Xi_{j,j'} = \langle \psi_j, \psi_{j'} \rangle \tag{10}$$

for an appropriately defined inner product $\langle \cdot, \cdot \rangle$. The Gram matrix of a set of vectors in an inner product space is the Hermitian matrix of inner products, whose entries are given by the inner product, $\langle \psi_j, \psi_{j'} \rangle$. In our context, ψ_j is a vector whose hth element is the impulse response of the target variable to shock j at the hth horizon in some set of horizons $\mathfrak{H}$. The matrix Ξ will thus depend on the reduced form parameters, $\{\{B_\ell\}_{\ell=1}^L, \Sigma\}$, as well as $\mathcal{H}$ in the time domain and Ω in the frequency domain, as we explain in Section 4. The solution for θ is the principal eigenvector (i.e., the eigenvector associated with the largest eigenvalue) of Ξ. We explain in the Online Appendix that it is rare for Ξ to have repeated non-zero eigenvalues and therefore (9) will generally have a unique solution.

Suppressing the dependence on $\{\{B_\ell\}_{\ell=1}^L, \Sigma\}$ for convenience, we have:

$$\Xi^{time}(\mathcal{H}) \equiv \sum_{h \in \mathcal{H}} \Psi_h' \delta_i \delta_i' \Psi_h, \tag{11}$$

$$\langle \psi_j, \psi_{j'} \rangle^{time} \equiv \psi_j \cdot \psi_{j'}, \tag{12}$$

from (4) in the time domain and:

$$\Xi^{freq}(\Omega) \equiv \text{Re}\left[\int_{\omega \in \Omega} \Upsilon(\omega)\, d\omega \right], \tag{13}$$

$$\langle \psi_j, \psi_{j'} \rangle^{freq} \equiv \int_{\omega \in \Omega} \Gamma_{i,j}^{\text{Re}}(\omega)\Gamma_{i,j'}^{\text{Re}}(\omega) + \Gamma_{i,j}^{\text{Im}}(\omega)\Gamma_{i,j'}^{\text{Im}}(\omega) d\omega, \tag{14}$$

from (6) in the frequency domain, where $\Gamma_{i,j}^{\text{Re}}(\omega)$ and $\Gamma_{i,j}^{\text{Im}}(\omega)$ are the real and imaginary parts of the (i,j) element of the transfer function, $\Gamma(\omega)$.[6] For the time domain, $\mathfrak{H} = \mathcal{H}$ and the dot product is clearly an inner product. In the frequency domain, $\mathfrak{H} = \mathbb{Z}_{\geq 0}$, the set of non-negative integers; we verify that $\langle \cdot, \cdot \rangle^{freq}$ defines an inner product in the Online Appendix.

As a Gram (and thus Hermitian) matrix, Ξ is analogous to a covariance matrix,

[6] Barsky and Sims (2011) and Ben Zeev and Khan (2015) replace (11) with $\sum_{k=1}^K \sum_{h=0}^{H_k} \Psi_h' \delta_i \delta_i' \Psi_h$, thus accumulating the FEV over a set of horizons $\{H_k\}_{k=1}^K$. Our general theorem will still apply, but the interpretation will have to account for the sum over $\{H_k\}_{k=1}^K$.

but of impulse responses and not variables.[7] Hence, even if the structural shocks are orthogonal, the off-diagonal elements of Ξ will be large (relative to the diagonals) if the shocks have similar dynamic effects on the target variable i.

4 Necessary and Sufficient Conditions for Validity

We now derive conditions that a generic Gram matrix Ξ must satisfy for max-share identification to be valid. We then discuss their interpretation in the time and frequency domain max-share problems. For convenience, we will henceforth assume without loss of generality that the true target structural shock is ordered first. Valid identification corresponds to the principal eigenvector of Ξ (and solution to (9)) being $\theta = \delta_1$, so that all the weight is placed on the true target structural shock.

4.1 General Result

Denote the *direct sum* of two square matrices X_1 and X_2 by $\mathsf{X}_1 \oplus \mathsf{X}_2$, i.e.,

$$\mathsf{X}_1 \oplus \mathsf{X}_2 \equiv \begin{bmatrix} \mathsf{X}_1 & 0 \\ 0 & \mathsf{X}_2 \end{bmatrix}.$$

In addition, let $\lambda_{\max}(\mathsf{X}_2)$ denote the largest eigenvalue of X_2 and let $\|\cdot\|$ denote the Euclidean norm for vectors or the Frobenius norm for matrices.

Theorem 1. *The unique solution to the general max-share problem:*

$$\arg\max_{\theta} \theta'\Xi\theta \text{ subject to } \theta'\theta = 1$$

is δ_1, i.e., max-share identification is valid, if and only if

- *(Orthogonality) $\Xi = \Xi_{1,1} \oplus \Xi_{2:N,2:N}$; and*
- *(Relative size) $\Xi_{1,1} > \lambda_{\max}(\Xi_{2:N,2:N})$.*

Theorem 1 is a key result that provides necessary and sufficient conditions for max-share identification to obtain the true target shock. It clarifies the economic

[7] It would be an actual covariance matrix if ψ_j and $\psi_{j'}$ were replaced with vectors each containing realizations of a random variable and the matrix was normalized by the number of observations.

reasoning required to justify the use of max-share, similar to how one would approach other identification schemes such as zero restrictions, sign restrictions, or instrumental variables. Concretely, for max-share to place all its weight on the first shock, two conditions must be satisfied.

First, the *orthogonality* condition states that Ξ must have the block diagonal structure:

$$\Xi = \begin{bmatrix} \Xi_{1,1} & 0 \\ 0 & \Xi_{2:N,2:N} \end{bmatrix}.$$

Interpreting Ξ as a covariance matrix, this corresponds to orthogonality between the first component and each of the other components. It restricts the shape of the impulse response of the target variable to the target shock relative to each of the other shocks. In the stylized example in Figure 1, Panel B showed that a seemingly small modification in the Shock 2 response (increasing the horizon 1 response from 0 to 0.2) in violation of orthogonality resulted in max-share placing over 40% of the weight on Shock 2, illustrating that the condition is not only stringent but also consequential. More generally, we will show that it is violated for a large class of impulse responses, challenging the understanding of many empirical applications of max-share.

As an empirically relevant example, consider the identification of TFP shocks by targeting output at a medium horizon (Uhlig, 2004a; Francis et al., 2014). The identification conditions are likely to be violated since the response of output to TFP shocks shares some similarity to its response to other shocks such as demand or borrowing cost shocks. An exception for which the shape becomes irrelevant is the case with $\Xi_{2:N,2:N} = 0$, i.e., all other shocks are fully dominated by the target shock over the target horizons or frequencies. This is unlikely in most settings besides two cases—long-run identification ($H \to \infty$) and internal instruments ($H = 0$).

Next, the *relative size* condition states that $\Xi_{1,1}$ must be larger than the largest eigenvalue of the lower block $\Xi_{2:N,2:N}$. Importantly, it is insufficient for $\Xi_{1,1}$ to be the largest diagonal element of Ξ. In particular, if the off-diagonal elements of $\Xi_{2:N,2:N}$ are large, then the principal eigenvalue of $\Xi_{2:N,2:N}$ can be larger than any of the individual diagonal elements. In the extreme case, when $\Xi_{2:N,2:N}$ is rank one, its principal eigenvalue will be its trace, i.e., $\lambda_{\max}(\Xi_{2:N,2:N}) = \sum_{j=2}^{N} \Xi_{j,j}$. The Panel D of Figure 1 presented a stylized example in which $\Xi_{1,1}$ was the largest diagonal element (since the response to Shock 1 was larger than that to Shocks 2 and 3) but the principal eigenvector had zero loading on the target Shock 1.

Both conditions emphasize the neglected role of impulse response shapes in the existing max-share literature. These shape restrictions apply to the responses of the target variable to the full set of shocks jointly. In the context of the VAR in (1) with iid shocks, this will depend on the internal propagation of the system through $\{B_\ell\}_{\ell=1}^L$, placing a potentially heavy burden of a priori knowledge on the econometrician. In contrast, the literature has tended to make broader claims without the caveats of Theorem 1, calling into question the validity of these identification strategies and their economic interpretation.[8]

4.2 Time Domain

Interpreting $\Xi^{time}(\mathcal{H})$. To apply Theorem 1 to the time domain, denote the response of the target variable to shock j over horizons $\mathcal{H}$ by vector $\psi_{\mathcal{H},j}$, so that the hth element of $\psi_{\mathcal{H},j}$ corresponds to response at the hth horizon in $\mathcal{H}$. We can then write the (j, j') entry of $\Xi^{time}(\mathcal{H})$ as follows:

$$\Xi^{time}_{j,j'}(\mathcal{H}) = \psi_{\mathcal{H},j} \cdot \psi_{\mathcal{H},j'} = \begin{cases} \|\psi_{\mathcal{H},j}\|^2 = \sum_{h\in\mathcal{H}} \Psi^2_{h,ij} & \text{if } j = j' \\ \|\psi_{\mathcal{H},j}\| \, \|\psi_{\mathcal{H},j'}\| \cos\alpha_{jj'}(\mathcal{H}) & \text{if } j \neq j' \end{cases}. \tag{15}$$

$\Xi^{time}(\mathcal{H})$ has the form of a covariance matrix, where $\|\psi_{\mathcal{H},j}\|^2$ corresponds to the jth variance and $\cos\alpha_{jj'}(\mathcal{H})$ plays the role of the correlation between impulse responses j and j'. The diagonal elements ($j = j'$) capture the squared magnitude of the response to the jth shock, $\|\psi_{\mathcal{H},j}\|^2$. The off-diagonal ($j \neq j'$) terms depend not only on the magnitude of the responses to j and j', but also to $\cos\alpha_{jj'}(\mathcal{H})$, where $\alpha_{jj'}$ is the angle between the vectors $\psi_{\mathcal{H},j}$ and $\psi_{\mathcal{H},j'}$. The ratio, $\Xi^{time}_{j,j'}(\mathcal{H})/\Xi^{time}_{j,j}(\mathcal{H})$, of an off-diagonal element with the diagonal element from that row is also the regression coefficient of $\psi_{\mathcal{H},j'}$ on $\psi_{\mathcal{H},j}$ without an intercept. Consequently, $\Xi^{time}_{j,j'}$ scales with the underlying impulse responses, $\psi_{\mathcal{H},j'}$ and $\psi_{\mathcal{H},j}$, but is not invariant to translating either of these responses (e.g., replacing $\psi_{\mathcal{H},j'}$ with $\psi_{\mathcal{H},j'} + \mathsf{c}$).

Identification Conditions. The orthogonality condition in the time domain is:

$$\psi_{\mathcal{H},1} \cdot \psi_{\mathcal{H},j} = 0 \text{ for all } j \neq 1, \tag{16}$$

[8]For example, a typical argument for using max-share is that the shock of interest is the primary driver of the target variable (e.g., Fève and Guay, 2019; Kurmann and Sims, 2021).

which, with $\|\psi_{\mathcal{H},j}\| > 0$, implies $\cos\alpha_{1j}(\mathcal{H}) = 0$ for all $j \neq 1$. Importantly, orthogonality here is defined in the space $\mathbb{R}^{|\mathcal{H}|}$, where $|\mathcal{H}|$ is the cardinality of the set $\mathcal{H}$. It depends on how the responses move between negative and positive values over different horizons and not on the correlation across innovations.

The condition is easily violated. For instance, if the target shock and some other shock $j \neq 1$ produce impulse responses that are strictly positive over a finite set of horizons $\mathcal{H}$, then we have $\psi_{\mathcal{H},1} \cdot \psi_{\mathcal{H},j} > 0$. This occurs, for instance, when both responses have the same shape as the impulse response of a stationary AR(1) to an iid shock. Alternatively, consider the common application of distinguishing TFP news and surprise shocks. The TFP news shock literature typically assumes that the shock has a zero or small effect on TFP on impact, but that its response grows over the target horizons and is persistent. In contrast, the TFP surprise shock produces a relatively large initial response in TFP, but its effect decays over time. Even though the two responses look markedly different, they are both strictly positive and thus violate (16) for any finite set of horizons, $\mathcal{H}$.

As we increase the number of shocks, N, condition (16) potentially becomes even more problematic because it has to be satisfied for all other shocks in the system. This requires the econometrician to be willing to make statements about the shape of each of the corresponding impulse responses, which potentially requires labeling even the untargeted shocks. For example, suppose we knew that the responses to shocks 2 to N have stationary AR(1) shapes. For the response to the target shock to be orthogonal to all of these simultaneously, the responses $\psi_{\mathcal{H},j}$ for $j \neq 1$ must be identical up to scale. However, this threatens the relative size condition, which becomes $\Xi^{time}_{1,1} > \sum_{j=2}^{N} \Xi^{time}_{j,j}$, since $\sum_{j=2}^{N} \Xi^{time}_{j,j}$ grows with N.

The relative size condition in the time domain can be written as:

$$\|\psi_{\mathcal{H},1}\|^2 > \lambda_{\max}\left(\Xi^{time}_{2:N,2:N}(\mathcal{H})\right). \tag{17}$$

As discussed before, $\lambda_{\max}(\Xi^{time}_{2:N,2:N})$ depends on both magnitude and correlations of target variable responses to untargeted shocks $2, \ldots, N$. The typical justification of $\|\psi_{\mathcal{H},1}\| > \|\psi_{\mathcal{H},j}\|$ for any $j \neq 1$ is a necessary but not sufficient condition since $\lambda_{\max}(\Xi^{time}_{2:N,2:N})$ could be large with small but similarly shaped responses. Equation (15) makes clear that the measure of similarity is $\cos\alpha_{jj'}$.

An important property of Ξ^{time} reflected in (15) is that $\|\psi_{\mathcal{H},j}\|$ is defined as a

sum over horizons $\mathcal{H}$. In the usual max-share implementation targeting the FEV at horizon H, we have $\mathcal{H} = \{0, \ldots, H\}$, and $\|\psi_{\mathcal{H},j}\|$ depends not just the relative size of individual responses at horizon H but instead on the entire response from impact through horizon H.[9] With a finite target horizon H, the responses at short horizons continue to impact the FEV. Consequently, the untargeted innovations need to not only have transitory effects, but the corresponding impulse responses also need to be negligible at short horizons unless they satisfy the orthogonality condition.

To prevent the dependence on untargeted horizons, Dieppe et al. (2021) propose a so-called *non-accumulated max-share* in which they set $\mathcal{H} = \{H\}$. The drawback of such an approach is that it implies $\psi_{\mathcal{H},j} \cdot \psi_{\mathcal{H},j'} = \|\psi_{\mathcal{H},j}\| \, \|\psi_{\mathcal{H},j'}\|$ or $\cos \alpha_{jj'} = 1$ for all (j, j') pairs. In other words, Ξ^{time} becomes rank one. Appendix B.1 shows that this represents an extreme violation of the orthogonality condition and the weight on each true shock is proportional to the size of impulse responses to it. This is consistent with the mixed success Dieppe et al. (2021) find in Monte Carlo simulations comparing non-accumulated max-share to the traditional max-share approach.

4.3 Frequency Domain

Interpreting $\Xi^{freq}(\Omega)$. In the frequency domain, the interpretation of orthogonality and relative size differs from the time domain. To derive the analogous objects, denote the real and imaginary parts of the transfer function, $\Gamma(\omega)$, in (8) by:

$$\Gamma^{\text{Re}}(\omega) \equiv \sum_{h=0}^{\infty} \Psi_h \cos(\omega h) \text{ and } \Gamma^{\text{Im}}(\omega) \equiv -\sum_{h=0}^{\infty} \Psi_h \sin(\omega h).$$

For variable i and shock j and frequency ω, characterize $\Gamma_{ij}(\omega)$ by the so-called *gain*:

$$\kappa_{i,j}(\omega) \equiv \sqrt{\left(\Gamma_{ij}^{\text{Re}}(\omega)\right)^2 + \left(\Gamma_{ij}^{\text{Im}}(\omega)\right)^2} \tag{18}$$

and *phase*:

$$\varphi_{i,j}(\omega) \equiv \tan^{-1}\left[-\frac{\Gamma_{ij}^{\text{Im}}(\omega)}{\Gamma_{ij}^{\text{Re}}(\omega)}\right] \tag{19}$$

[9] See, for example, the main assumption in Francis et al. (2014) that all untargeted shocks have transitory effects on labor productivity.

so that $\Gamma_{ij}(\omega) = \kappa_{i,j}(\omega)e^{i\varphi_{i,j}(\omega)}$. The gain captures how much shock j is amplified in the target variable's frequency ω component. It increases proportionally with the standard deviation of shock j. The phase captures how much shock j is shifted back in time relative to target variable's frequency ω component. For example, the phase corresponding to Shock 1 in Figure 1 will be ω, as the impulse response shifts the shock back by $\varphi_{i,j}(\omega)/\omega = 1$ period.[10]

We can write the (j, j') entry of $\Xi^{freq}(\Omega)$ as:

$$\Xi^{freq}_{j,j'}(\Omega) = \begin{cases} \int_{\omega\in\Omega} \kappa^2_{i,j}(\omega)d\omega & \text{if } j = j' \\ \int_{\omega\in\Omega} \kappa_{i,j}(\omega)\kappa_{i,j'}(\omega)\cos\left(\varphi_{i,j}(\omega) - \varphi_{i,j'}(\omega)\right)d\omega & \text{if } j \neq j' \end{cases}, \tag{20}$$

which parallels the expression for $\Xi^{time}_{j,j'}$ in (15). For intuition, focus on the objects inside the integral, which are the elements of $\Xi^{freq}_{j,j'}(\Omega)$ when we consider a singleton frequency band, $\Omega = \{\omega\}$. The squared gain, $\kappa^2_{i,j}$ now takes the place of the squared norm, $\|\psi_{\mathcal{H},j}\|^2$, playing the analogous role of capturing the size of the impulse response at frequency ω. On the off-diagonal, the angle, $\alpha_{j,j'}$, between impulse responses is now replaced by the phase difference, $\varphi_{i,j}(\omega) - \varphi_{i,j}(\omega)$.

Identification Conditions. In the frequency domain, the orthogonality condition is:

$$\int_{\omega\in\Omega} \kappa_{1,1}(\omega)\kappa_{1,j'}(\omega)\cos\left(\varphi_{1,1}(\omega) - \varphi_{1,j'}(\omega)\right)d\omega = 0 \quad \text{for all} \quad j \neq 1, \tag{21}$$

and the relative size condition can be written:

$$\int_{\omega\in\Omega} \kappa^2_{1,1}(\omega)d\omega > \lambda_{\max}\left(\Xi^{freq}(\Omega)_{2:N,2:N}\right). \tag{22}$$

Even for a single frequency, ω, the elements of Ξ^{freq} depend on the impulse response over all horizons $h \geq 0$, subject to the weights $e^{-i\omega h}$. Since these weights are periodic, they do not discriminate between short and long horizons. While the frequency ω is often connected with $\tau = 2\pi/\omega$ periods (Stock and Watson, 1999), this association refers to the periodicity of $e^{-i\omega h}$ and does not imply some special relevance of the impulse response at horizon τ, as emphasized by Angeletos et al. (2020). More generally, (20) requires integrating nonlinear functions of the time-domain representation of the impulse responses over the frequency band, Ω. Consequently, it is arguably a

[10] See Watson (2001) for a more detailed overview and additional examples.

challenging task to justify the use of max-share identification with reference to specific assumptions on Ξ^{freq}.

The phase difference presents a particularly stark contrast with the time domain. Consider the two responses: $\psi_1 = (0, 1, 0, 0, \ldots)'$ and $\psi_2 = (1, 0, 0, \ldots)'$. While these are orthogonal by our time domain criteria for a given time horizon $\mathcal{H} = \{0, \ldots, H\}$, their phase difference in the frequency domain is ω. Therefore, in the frequency domain, they are only orthogonal at frequency $\pi/2$, which is typically associated with fluctuations at the 2-period frequency, despite the impulse response peaks differing only by 1 period.[11]

4.4 Comparison to Other Identification Schemes

We now contrast the max-share identification conditions with some common structural VAR identification approaches in order to provide additional intuition.

Internal Instruments. In the time domain, max-share with target horizon $H = 0$ is equivalent to using the target variable as an "internal instrument" (Noh, 2018; Plagborg-Møller and Wolf, 2021), i.e., ordering it first and computing a Cholesky decomposition of Σ so that C is lower triangular. This places zero restrictions on the $(1, j)$ elements on the top row of C for $j > 1$. From the lens of the max-share problem, setting $H = 0$ makes Ξ rank one. Consequently, the zero restrictions are necessary for orthogonality: with only one horizon in $\mathcal{H} = \{0\}$ and a non-zero response to the target shock, the orthogonality condition can only be satisfied if the responses to all untargeted shocks are zero on impact. Otherwise, the weight on each true shock is proportional to the size of the impulse response to it, as shown in Appendix B.1.

Targeting the FEV for horizon $H > 0$ instead, the restrictions no longer apply to the response on impact but to the entire response over horizons $0, \ldots, H$. The zero restrictions are replaced by the assumption that the remaining impulse responses lie in the null space of $\psi_{0:H,1}$, which is arguably harder to satisfy in most settings. Specifically, unless the untargeted responses happen to be orthogonal to the targeted one, we would require that the target variable does not respond to untargeted shocks both at horizon 0 and at all other horizons up to H.

[11] See Angeletos et al. (2020) for an empirical example comparing max-share in the time and frequency domains.

Sign Restrictions. While max-share does not directly impose sign restrictions, it can still indirectly imply joint conditions on the signs of the responses to the targeted and untargeted shocks. For instance, in the time domain orthogonality is violated if the responses of the target variable to the target shock and at least one untargeted shock have the same signs over the horizons $\mathcal{H}$. In contrast to identification via sign restrictions where the assumptions are on the response of only the target shock, the max-share identification conditions apply to both the targeted and untargeted shocks jointly. In other words, one needs to label the untargeted shocks and have a theory for them to justify max-share identification.

Instrumental Variables. The orthogonality condition in Theorem 1 brings to mind analogous exogeneity conditions for external instrument VARs or even single equation linear regressions. However, whereas exogeneity conditions in other settings pertain to the covariance between residuals and regressors (e.g., Sargan, 1958; Engle et al., 1983; Stock and Watson, 2018), in our context orthogonality is required between impulse responses. The former requires theoretical underpinnings for the sources of disturbances. The latter is dictated by the propagation of these disturbances. The two do not nest each other. For example, in an AR(1) process with two innovations that are iid over time, the two innovations will produce identical responses ($\alpha_{12} = \varphi_{i,1}(\omega) - \varphi_{i,2}(\omega) = 0$) regardless of their correlation with each other.

5 Beyond Exact Identification

In practice, it is infeasible to fully verify the conditions of Theorem 1 since we do not observe all the true structural shocks. However, it is possible to obtain a lower bound on the degree of contamination or weaken the identification conditions if the econometrician has additional knowledge.

5.1 Practical Diagnostics and Refinements

We first present a theorem that provides a practical check of the validity of max-share identification in empirical applications. Unlike the conceptual approach in the previous section, this implementation uses output that is directly available from the estimation.

Using Untargeted Responses to Bound Contamination. With knowledge about the response of the target variable to an untargeted shock, we can obtain a measure of how much the max-share identified shock is contaminated by this untargeted shock. The untargeted shock can be externally identified using any identification scheme as long as it is plausibly purged of the true target shock. For example, one could consider a shock identified via instrumental variables if the instrument is exogenous with respect to the targeted shock, even if it captures a combination of true untargeted shocks.

The following theorem gives practitioners a straightforward way to either confirm that the max-share shock is contaminated or rule out the likelihood of it being contaminated by the externally identified shock. Crucially, its implementation uses output directly available from the estimation.

Theorem 2. *Suppose the max-share problem (9) has a unique solution* $\theta = (\theta_1, \ldots, \theta_N)'$ *with the associated largest eigenvalue* $\lambda_{\max}(\Xi)$ *and max-share impulse response* $\psi^* = \sum_{k=1}^{N} \theta_k \psi_k$. *Then for an impulse response* $\hat{\psi} \equiv \sum_{j=2}^{N} \alpha_j \psi_j$ *with* $\sum_{j=2}^{N} \alpha_j^2 = 1$, *we have:*

$$\langle \psi^*, \hat{\psi} \rangle = \langle \psi^*, \psi^* \rangle \sum_{j=2}^{N} \alpha_j \theta_j = \lambda_{\max}(\Xi) \sum_{j=2}^{N} \alpha_j \theta_j. \tag{23}$$

Furthermore, we have the following upper bound for the weight on the targeted shock:

$$\theta_1^2 \leq 1 - \left(\frac{\langle \psi^*, \hat{\psi} \rangle}{\langle \psi^*, \psi^* \rangle} \right)^2. \tag{24}$$

Theorem 2 shows that even though the max-share problem itself does not allow us to directly observe the weights, θ, on the true structural shocks, we can indirectly learn about the contamination as long as we can identify the untargeted shocks that we are concerned about.[12] Suppose we observe Shock 2 (or its impulse response), i.e., $\alpha_2 = 1$ and $\alpha_j = 0$ for $j > 2$. Then (23) states that we can obtain θ_2 by projecting the impulse response to the observed Shock 2, ψ_2, on the impulse response of the target

[12]The impulse response to the max-share identified shock will be orthogonal to the implied set of untargeted shocks by construction. However, these implied shocks are not the true structural shocks. This is analogous to the fitted residuals in an ordinary least squares regression satisfying the orthogonality conditions by construction even in the presence of endogeneity.

variable to the max-share shock, ψ^*. The bound (24) then follows, giving an upper bound on the weight that max-share identification is placing on the target shock.[13]

We provide numerical and empirical examples below to show that the the magnitude of $\frac{\langle \psi^*, \hat{\psi} \rangle}{\langle \psi^*, \psi^* \rangle}$ can be substantial, bounding θ_1 considerably away from 1. Moreover, with $N > 2$, the upper bound (24) will not be attained unless the remaining untargeted shocks produce impulse responses that are orthogonal to the targeted one. As we have argued, this is unlikely in practice. Therefore, arguing that the target variable is driven by a single shock or modeling a single shock to match the impulse responses of various variables to the max-share shock (e.g., Angeletos et al., 2020; Basu et al., 2025) based on a large FEV contribution can be misleading. Max-share only provides an upper bound for the contribution of the true targeted shock (by construction from the problem (9) and pointed out by Fève and Guay (2019)); Theorem 2 emphasizes that this could appreciably overstate the targeted shock's importance.

For concreteness, suppose we identified a TFP news shock using max-share with TFP as the target variable. The shock should be orthogonal to monetary policy shocks. Therefore, one can separately identify a monetary policy shock using an external instrument based on high frequency identification (Gertler and Karadi, 2015; Bauer and Swanson, 2023a) and obtain the corresponding impulse response of TFP, which plays the role of $\hat{\psi}$ in this example. As long as the identified monetary policy shock is itself not contaminated by TFP news shocks, the assumption that $\hat{\psi}$ places zero weight on ψ_1 is satisfied. Theorem 2 suggests projecting the impulse response of TFP to the monetary policy shock on the response to the max-share shock. A large projection coefficient is evidence that the max-share shock is placing substantial weight on the monetary policy shock. A small coefficient shows that the max-share shock is not materially contaminated by the identified monetary shock, but does not rule out potential contamination by other untargeted shocks.[14]

Even without directly observing or identifying the untargeted shock, we can utilize (23) and (24) in Theorem 2 to quantify the implied contamination of the max-share shock as long as we have information about the untargeted impulse response $\hat{\psi}$. Such

[13] Since θ_j is also the correlation between the max-share shock and Shock j in population, Theorem 2 provides an alternative to obtaining the correlation between the max-share shock and other shocks, but without sampling error (under the maintained assumption that we know the reduced form parameters).

[14] Miranda-Agrippino and Ricco (2021) and Bauer and Swanson (2023b) further discuss the information contained in high frequency monetary policy instruments. Meier and Reinelt (2024) show how TFP can respond to monetary policy shocks.

information can be obtained using the approach proposed by Plagborg-Møller (2019) to form priors on impulse responses. If the information is obtained from a structural model (Ingram and Whiteman, 1994; Del Negro and Schorfheide, 2004), we can also determine how the structural parameters determine the degree of contamination.

Max-Share with Constraints. To overcome concerns over identification, the literature has proposed extensions, most prominently incorporating zero restrictions (Barsky and Sims, 2011) or controlling for the observed shock before implementing max-share (Cascaldi-Garcia and Galvao, 2021; Basu et al., 2025).

Both examples above can be expressed as linear constraints in the following constrained max-share problem:

$$\arg\max_{\theta} \theta'\Xi\theta \text{ subject to } K'\theta = 0 \text{ and } \theta'\theta = 1, \tag{25}$$

for a full rank $K \in \mathbb{R}^{N\times m}$ with $1 \leq m \leq N-1$. The assumption that the target variable does not respond to the target shock on impact corresponds to $K = \Psi_0'\delta_1$.[15] Controlling for some observed shock j corresponds to $K = \delta_j$.

The theorem below shows that the constrained max-share problem (25) has identification conditions analogous to Theorem 1. Online Appendix B.2 presents a theorem that similarly parallels Theorem 2, but for the constrained problem.

Theorem 3. *The unique solution to the general constrained max-share problem* (25) *is* δ_1*, i.e., max-share identification is valid under the* m *linear constraints characterized by* K*, if and only if*

- *(Orthogonality)* $\check{\Xi} = \check{\Xi}_{1,1} \oplus \check{\Xi}_{2:N,2:N}$*;*
- *(Relative size)* $\check{\Xi}_{1,1} > \lambda_{\max}(\check{\Xi}_{2:N,2:N})$*; and*
- *(Feasibility)* $K'\delta_1 = 0$*,*

where $\check{\Xi} \equiv M_K \Xi M_K$ *and* $M_K \equiv I - K(K'K)^{-1}K'$ *is the annihilator matrix of* K*.*

The identification conditions for the constrained problem (25) in Theorem 3 are similar to those for the unconstrained problem (9) in Theorem 1 except for two

[15] To see this, note that the on-impact response of the max-share shock is $\Psi_0\theta$. The zero restriction requires its first element be zero, $\delta_1'\Psi_0\theta = 0$.

differences. First, there is an additional *feasibility* constraint to ensure that δ_1 is not ruled out as a solution by the constraint. Second, the matrix Ξ is replaced with $\check{\Xi}$, which is also a Gram matrix but with:

$$\check{\Xi}_{jj'} = \langle \sum_{k=1}^{N} M_{K,kj}\psi_k, \sum_{k=1}^{N} M_{K,kj'}\psi_k \rangle.$$

In other words, the vectors ψ_j underlying Ξ are replaced by $\sum_{k=1}^{N} M_{K,kj}\psi_k$, which capture the impulse responses of the target variable to linear combinations of shocks. The annihilator matrix, M_K, ensures that these linear combinations satisfy the constraint by projecting the impulse responses onto the subspace that is orthogonal to K, similar to its role in generating residuals in a linear regression context. Feasibility ensures that $\sum_{k=1}^{N} M_{K,k1}\psi_k = \psi_1$, so that the first row and column of $\check{\Xi}$ correspond to the first shock as they do for Ξ.

As long as feasibility is satisfied, the additional constraints $K'\theta = 0$ help to weaken the identification conditions. In the unconstrained problem (9) we require that ψ_1 be orthogonal to the space spanned by $\psi_2, \ldots, \psi_N$. In the constrained problem (25) each additional constraint (i.e., column of K) reduces the dimension of that space by 1. The reduced dimension also implies that $\lambda_{\max}(\check{\Xi}_{2:N,2:N}) \leq \lambda_{\max}(\Xi_{2:N,2:N})$. This is formalized in the following corollary.

Corollary 1. *Consider the constrained max-share problem (25). If the feasibility condition, $K'\delta_1 = 0$, holds, then the orthogonality and relative size conditions in Theorem 1 are sufficient but not necessary for those in Theorem 3:*

- $\Xi = \Xi_{1,1} \oplus \Xi_{2:N,2:N} \implies \check{\Xi} = \check{\Xi}_{1,1} \oplus \check{\Xi}_{2:N,2:N}$*; and*
- $\Xi_{1,1} > \lambda_{\max}(\Xi_{2:N,2:N}) \implies \check{\Xi}_{1,1} > \lambda_{\max}(\check{\Xi}_{2:N,2:N})$.

Nevertheless, the difficulties in exactly satisfying the identification conditions remain unless the constraint fully controls for the contaminating shocks. This is only possible if the number of constraints matches the number of untargeted shocks impacting the target variable. Moreover, the number of variables, N, limits the number of constraints. For example, we can control for at most $N-1$ observed shocks. Consequently, one is not at liberty to add an arbitrarily large number of "controls" before utilizing max-share (in contrast to a linear regression where we can continually add

regressors, sample size permitting) even though it is not uncommon to have numerous versions of a single class of shocks, as exemplified, for instance, in the contrasting approaches to identifying monetary policy shocks (McKay and Wolf, 2023; Brennan, Jacobson, Matthes, and Walker, 2024).[16] The constraints generally increase the share of the weight placed on the target shock, $\theta_1 / \sum_{j=1}^{N} |\theta_j|$, but it is possible to construct counterexamples in which this is not true.

5.2 Perturbations and Global Bounds

Besides two special cases—with Ξ rank one and with $N = 2$—that we discuss in Online Appendix B.1, closed form solutions are not generally available for the max-share problem (9) when the identification conditions in Theorem 1 are not satisfied. Nevertheless, we can obtain local approximations for small deviations from orthogonality and bounds for the solution with larger deviations.

Local Deviations from Orthogonality. The following proposition characterizes the solution when deviations from the orthogonality condition are small.

Proposition 1. *Suppose* $\Xi := \Xi_{1,1} \oplus \Xi_{2:N,2:N} + \mathrm{d}\Xi$, *where* $\Xi_{1,1} > \lambda_{\max}(\Xi_{2:N,2:N})$, *and*

$$\mathrm{d}\Xi = \begin{bmatrix} 0 & \nu' \\ \nu & \mathbf{0}_{(N-1)\times(N-1)} \end{bmatrix} \text{ with } \|\nu\| = o\left(\sqrt{\Xi_{1,1}^2 + \|\Xi_{2:N,2:N}\|^2}\right)$$

and $\lambda_{\max}(\Xi) = \Xi_{1,1} + O\left(\|\nu\|^2\right)$ *is simple, then the max-share problem has solution*

$$\left[1, \sum_{j=2}^{N} \frac{w_{0j}'\nu}{\Xi_{1,1} - \lambda_{0j}} w_{0j}'\right]' + O(\|\nu\|^2), \tag{26}$$

where $\{w_{0j}\}_{j=2}^{N}$ *is a complete set of orthonormal eigenvectors of* $\Xi_{2:N,2:N}$ *corresponding to (possibly repeated) eigenvalues* $\{\lambda_{0j}\}_{j=2}^{N}$.

Proposition 1 describes how the principal eigenvector changes with a particular perturbation in Ξ. The perturbation, $\mathrm{d}\Xi$, is zero on the diagonal blocks and characterized by the vector ν off the diagonal blocks. Intuitively, we can think of this

[16] See also the discussion of invertibility in VARs in Section 4.1.2 of Stock and Watson (2016) and references therein.

as leaving the responses to the untargeted shocks unchanged while perturbing the response to the targeted shock in such a way that its magnitude is unchanged but its shape is closer to the other responses.

The resulting expression, (26), has a straightforward interpretation. The first element of the principal eigenvector is unchanged up to first order. But the rest of the vector now reflects a shift in the direction of each of the eigenvectors, $\{w_{0j}\}_{j=2}^{N}$, of $\Xi_{2:N,2:N}$. The weight on each of these eigenvectors depends on the coefficients, $\frac{w_{0j}'\nu}{\Xi_{1,1}-\lambda_{0j}}$. The numerator is a projection, $w_{0j}'\nu$, of the perturbation, ν, on the eigenvector, w_{0j}. Thus, the principal eigenvector will shift more in the direction of w_{0j} if the perturbation of the targeted response brings it closer to that eigenvector. The denominator implies that the change in the principal eigenvector depends on the size of the eigenvalues λ_{0j} relative to $\Xi_{1,1}$. In other words, if the target impulse response is much larger than the other responses so that $\Xi_{1,1}$ is large relative to λ_{0j}, then small deviations from orthogonality will have a relatively minor effect on identification.

Global Bounds. With larger deviations, we can derive bounds that apply globally.

Proposition 2. *Suppose* $\Xi := \Xi_{1,1} \oplus \Xi_{2:N,2:N} + \mathrm{d}\Xi$, *where* $\Xi_{1,1} > \lambda_{\max}(\Xi_{2:N,2:N})$, *and*

$$\mathrm{d}\Xi = \begin{bmatrix} 0 & \nu' \\ \nu & \mathbf{0}_{(N-1)\times(N-1)} \end{bmatrix}$$

and $\lambda_{\max}(\Xi)$ *is simple, then the unique solution* θ *to the max-share problem satisfies:*

$$\sin\alpha(\delta_1,\theta) \leq \frac{2\,\|\nu\|}{\Xi_{1,1}-\lambda_{\max}(\Xi_{2:N,2:N})} \;\; \textit{and} \;\; \|\theta-\delta_1\| \leq \frac{2^{3/2}\,\|\nu\|}{\Xi_{1,1}-\lambda_{\max}(\Xi_{2:N,2:N})}, \qquad (27)$$

where the $\alpha(\delta_1,\theta)$ *is the principal angle between* δ_1 *and* θ, *and the sign of* θ_1 *is normalized to be positive.*

Proposition 2 applies the same perturbation, $\mathrm{d}\Xi$, as Proposition 1, but no longer requires that $\|\nu\|$ is small. The two expressions in (27) give alternative ways to measure the difference between δ_1 and θ, either through the principal angle, $\alpha(\cdot,\cdot)$, or the norm, $\|\cdot\|$, of the difference. In both cases, the denominator of the bound is the difference, $\Xi_{1,1}-\lambda_{\max}(\Xi_{2:N,2:N})$, between the two largest eigenvalues of Ξ. If the difference is sufficiently large, then even nontrivial deviations from orthogonality will

result in small degrees of contamination (an extreme case being exact identification when $\Xi_{2:N,2:N} = 0$, as mentioned in Section 4.1).[17]

6 Numerical Examples

We now consider a series of numerical examples that all follow the general form for a log-linearized dynamic stochastic general equilibrium (DSGE) model (Fernández-Villaverde, Rubio-Ramírez, Sargent, and Watson, 2007):

$$x_t = Fx_{t-1} + Q\varepsilon_t \tag{28}$$

$$y_t = Gx_{t-1} + R\varepsilon_t, \tag{29}$$

where $\varepsilon_t \sim \mathcal{N}(0, I)$. The impulse response of a variable $y_{i,t}$ to shock j at horizon h is:

$$\Psi_{i,j,h} = \begin{cases} R_{i,j} & h = 0 \\ G_{i,:}F^{h-1}Q_{:,j} & h \geq 1 \end{cases}. \tag{30}$$

In what follows, we abstract from whether the VAR can produce these responses and focus more abstractly on what they imply for Ξ and the identification conditions.[18]

Differences in the dynamic responses to different shocks beyond horizon 0 depend on differences in exposures, $Q_{:,j}$, of the states x_t to the shocks and dynamics of those states, as captured by F. If the states have similar autoregressive properties or exposures to the structural shocks, then a given variable will have a similar dynamic response to different shocks and max-share will result in a convolution of these shocks.

For intuition, consider the case where x_t and thus F are scalars. An example is the real business cycle model from King, Plosser, and Rebelo (1988) with iid shocks, where the only state, x_t, is (percentage deviations of) capital, $\hat{k}_t$. The response of $y_{i,t}$ to any shock then resembles an ARMA(1,1) with the same AR coefficient F. Suppose

[17] Consider a Blanchard and Quah (1989) economy where only one shock has a permanent effect on the target variable. With a sufficiently long horizon, one can get a large weight on the shock of interest even if orthogonality is not closely approximated. For instance, consider responses $\psi_{1,h} = 1 - 0.9^h$ and $\psi_{2,h} = 0.9^h$. For $\mathcal{H} = \{0, \ldots, 100\}$, we then have $\|\psi_{1,\mathcal{H}}\|^2 / \|\psi_{2,\mathcal{H}}\|^2 = 16$ and max-share places 95 percent of the weight on Shock 1 despite orthogonality being violated with $\psi_{1,\mathcal{H}} \cdot \psi_{2,\mathcal{H}} / \|\psi_{1,\mathcal{H}}\| \|\psi_{2,\mathcal{H}}\| = 0.25$. Nonetheless, there are other drawbacks to long-run identification arising from low frequency fluctuations (Francis and Ramey, 2009; Gospodinov et al., 2013).

[18] See Fernández-Villaverde et al. (2007) for conditions under which a VAR can reproduce the model impulse responses.

as before that the target shock is ordered first. Then, for a given horizon or frequency band and $R_{i,1}$, orthogonality requires a specific $R_{i,j}/Q_{i,j}$ for all $j \neq 1$.

6.1 ARMA(1,1) Impulse Responses

When impulse responses resemble those of an ARMA model, orthogonality is possible but highly dependent on parameters. As an example, we study the orthogonality condition and deviations from it in (28)-(29) with $F = \text{diag}\{\rho_j\}$, $Q = \text{diag}\{\rho_j - \phi_j\}$ with $|\rho_j| < 1$ for $j \in \{1, 2\}$, and $G = R = (1, 1)'$. The impulse responses then correspond to those of an ARMA(1,1) with the MA representation $(1 - \phi_j L)(1 - \rho_j L)^{-1}\varepsilon_{j,t}$, where L is the lag operator and ϕ_j and ρ_j are the MA and AR coefficients, respectively. The impulse responses are:

$$\psi_{j,h} = \begin{cases} 1 & h = 0 \\ (\rho_j - \phi_j)\rho_j^{h-1} & h > 0 \end{cases}. \tag{31}$$

Time Domain. Taking $\mathcal{H} = \{0, \ldots, H\}$ as is common in the time domain, we have:

$$\psi_{1,\mathcal{H}} \cdot \psi_{2,\mathcal{H}} = 1 + (\rho_1 - \rho_1)(\rho_2 - \phi_2)\frac{1 - (\rho_1\rho_2)^H}{1 - \rho_1\rho_2} \tag{32}$$

and orthogonality is achieved when:

$$H = \log\left(1 + \frac{1 - \rho_1\rho_2}{(\rho_1 - \phi_1)(\rho_2 - \phi_2)}\right) / \log(\rho_1\rho_2). \tag{33}$$

The solution to H and deviations from orthogonality depend on the parameter values. For instance, with $\phi_2 = 0$, we only have a positive solution for H if $\phi_1\rho_2 > 1$. Moreover, the solution to H may not even be an integer.

As an illustration, we set $\rho_1 = 0.75$, $\rho_2 = 0.95$, $\phi_2 = 0$, and vary $\phi_1\rho_2 \in \{0, 1, 1.5\}$. Figure 2 shows that with $\phi_1 = 0$, both responses have AR(1) shapes. The persistence of each impulse response is sufficiently distinct to make the responses easily distinguishable visually. Nevertheless, we have severe violations of orthogonality, with $\cos\alpha_{12}$ decreasing with H from 1 at $H = 0$ to 0.72 as $H \to \infty$. In contrast, with $\phi_1\rho_2 = 1$, orthogonality is approximated for all but the shortest horizons and attained when we take $H \to \infty$. When we increase ϕ_1 further such that $\phi_1\rho_2 = 1.5$, we again have relatively strong violations of orthogonality except around the solution to (33) of

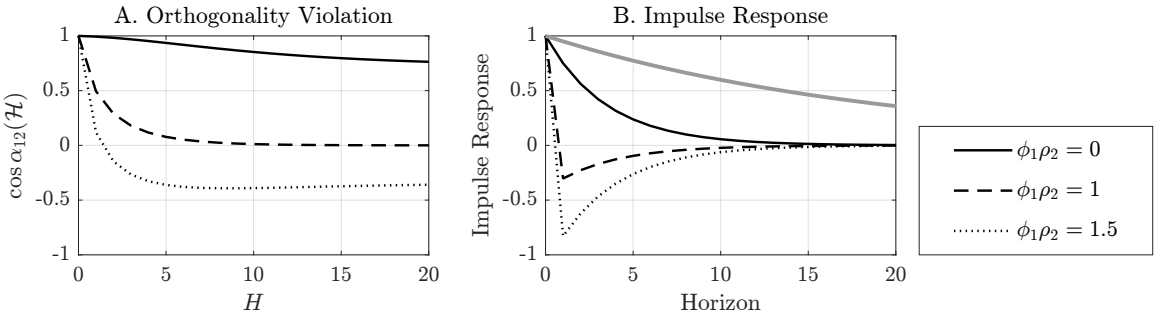


Figure 2: Orthogonality violations in the time domain for ARMA(1,1) impulse responses. **Panel A:** Orthogonality violations, $\cos \alpha_{12}(\mathcal{H})$; **Panel B:** Impulse responses. Black lines correspond to Shock 1; gray line corresponds to Shock 2.

2.8. The contrast is striking—a theory justifying the $\phi_1\rho_2 = 1$ case is unlikely to rule out $\phi_1\rho_2 = 1.5$ given how similar the responses are, but the two parameterizations have very different implications for the orthogonality condition. Therefore, while one can construct examples in which max-share approximately or exactly identifies the true shock of interest, the performance is very sensitive to the details of the responses.

Frequency Domain. As we did in the time domain, we now study the orthogonality conditions in the class of impulse responses defined in (31). To simplify our analysis, we focus on the single frequency case with $\Omega = \{\omega\}$.

Defining:

$$\varphi_j^{\mathsf{x}}(\omega) = -\tan^{-1}\left[\frac{\mathsf{x}_j \sin\omega}{1 - \mathsf{x}_j \cos\omega}\right] \quad \text{for} \quad \mathsf{x} \in \{\rho, \phi\}, \tag{34}$$

we can write the phase difference as:

$$\varphi_1(\omega) - \varphi_2(\omega) = (\varphi_1^{\rho}(\omega) - \varphi_1^{\phi}(\omega)) - (\varphi_2^{\rho}(\omega) - \varphi_2^{\phi}(\omega)). \tag{35}$$

With $\phi_1 = \phi_2 = 0$, i.e., AR(1) impulse responses, we can show that $\cos(\varphi_1 - \varphi_2) > 0$ and orthogonality is not satisfied for any frequency ω as in the time domain. Equation (35) shows that the MA component in response j induces a phase shift by $\varphi_j^{\phi}(\omega)$ that can, for the right parameterization and frequency, yield orthogonality.

Similar to the time domain, we set $\rho_1 = 0.75$, $\rho_2 = 0.95$, and $\phi_2 = 0$ as an illustration. We now vary $\phi_1 \in \{-1.5, 0, 1.5, 2\}$. With $\phi = 0$, we have large deviations from orthogonality, with $\cos(\varphi_1(\omega) - \varphi_2(\omega)) \geq 0.72$. More broadly, there is no $\omega \in (0, \pi)$ satisfying orthogonality for $\phi_1 \in (-1, 0.86)$. With other values of ϕ_1, we do

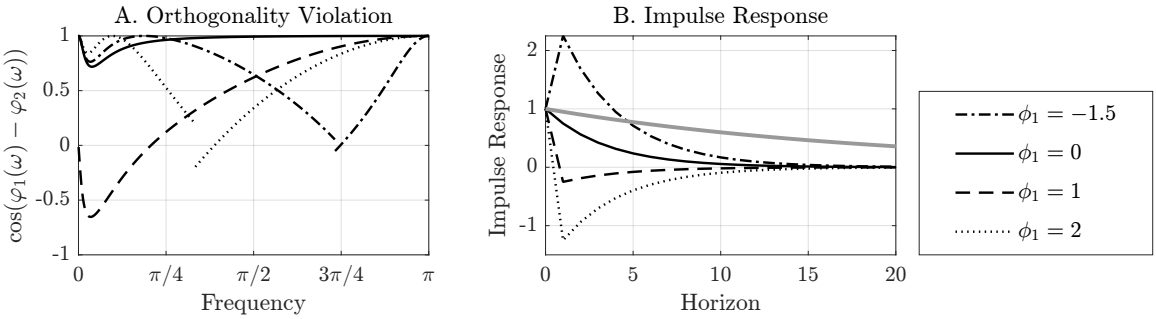


Figure 3: Orthogonality violations in the frequency domain for ARMA(1,1) impulse responses. **Panel A:** Orthogonality violations, $\cos(\varphi_1(\omega)-\varphi_2(\omega))$; **Panel B:** Impulse responses. Black lines correspond to Shock 1; gray line corresponds to Shock 2.

find frequencies for which orthogonality is satisfied—$\omega = 2.4$, $\omega = 0.6$, and $\omega = 1.2$ for $\phi_1 = -1.5$, $\phi_1 = 1$, and $\phi_1 = 2$, respectively. However, the solutions are sensitive to and non-monotonic in ϕ_1. In addition, the curves for $\cos(\varphi_1(\omega)-\varphi_2(\omega))$ on Panel A of Figure 3 are highly nonlinear and even discontinuous. Finally, around the values of ω satisfying orthogonality, $\cos(\varphi_1(\omega)-\varphi_2(\omega))$ has a relatively steep gradient. These observations raise doubt about the plausibility of max-share identification in the frequency domain as it requires precise beliefs on the shape of the impulse responses in order to satisfy the orthogonality condition for the frequency of interest.

6.2 Supply and Demand

Another instructive case takes F and Q to be diagonal and $R = GQ$, so that x_t consists of independent AR(1) processes and the model can be written as:

$$y_t = Gx_t = GFG^{-1}y_{t-1} + R\varepsilon_t. \tag{36}$$

Here, we consider the particular case of a supply and demand system:

$$q_t^s = \gamma^s p_t + \eta_t^s, \tag{37}$$

$$q_t^d = -\gamma^d p_t + \eta_t^d, \tag{38}$$

where q_t^s, q_t^d, and p_t denote, respectively, the quantity supplied, quantity demanded, and price in logs. We assume the shocks follow AR(1) processes:

$$\eta_t^{\mathsf{x}} = \rho^{\mathsf{x}} \eta_{t-1}^{\mathsf{x}} + \sigma^{\mathsf{x}} \varepsilon_t^{\mathsf{x}}, \tag{39}$$

where $\varepsilon_t^{\mathsf{x}} \overset{iid}{\sim} \mathcal{N}(0,1)$ with $\mathsf{x} \in \{s, d\}$. Using the fact that $q_t = q_t^s = q_t^d$ in equilibrium, we can express $y_t \equiv (q_t, p_t)'$ in the form (36). The impulse responses will propagate as AR(1) processes with persistence ρ^s and ρ^d, inheriting the dynamics of the shocks since the rest of the model (37)–(38) is static.[19] Details are in the Online Appendix.

We identify the supply shock using max-share identification on output in the time domain for a large but finite H. The approach follows Uhlig (2004a) and Francis et al. (2014), who use max-share as an alternative for long-run identification (targeting labor productivity instead of output). We consider the parameter values:

$$\begin{aligned} \gamma^s &= 1.00, \qquad \rho^s = 1.00, \qquad \sigma^s = 1.00 \\ \gamma^d &= 0.50, \qquad \rho^d = 0.95, \qquad \sigma^d = 1.50 \end{aligned}$$

The supply shock is a random walk and the only shock with a permanent effect on q_t. The demand shock has persistence, $\rho^d = 0.95$, less than the value of 0.98 used by Francis et al. (2014). Following Francis et al. (2014), we set $\mathcal{H} = \{0, \ldots, 40\}$.

Figure 4 shows the impulse responses to the true shocks as well as those obtained using max-share identification. While the impulse response of output to the true demand shock is about three times the size of the supply shock on impact, this is reversed by horizon $h = 40$, with the impulse response to the supply shock now three times that of the demand shock. Despite the relatively small impulse response to the demand shock at horizon 40, Panel C in Figure 4 shows that the true supply shock only accounts for less than a third of the FEV, suggesting that max-share identification is unlikely to perform well.

The max-share shock differs substantially from the true supply shock. First, Panels A and B show that the max-share shock produces a positive response in both output and price, resembling a demand shock. Next, Panel C shows that the max-share shock has a contribution of close to one, roughly three times the FEV contribution of the true supply shock. These results arise because the max-share

[19] The three-equation New Keynesian model (An and Schorfheide, 2007; Galí, 2015) has a similar structure in which persistence arises solely through the dynamics of the exogenous shocks.

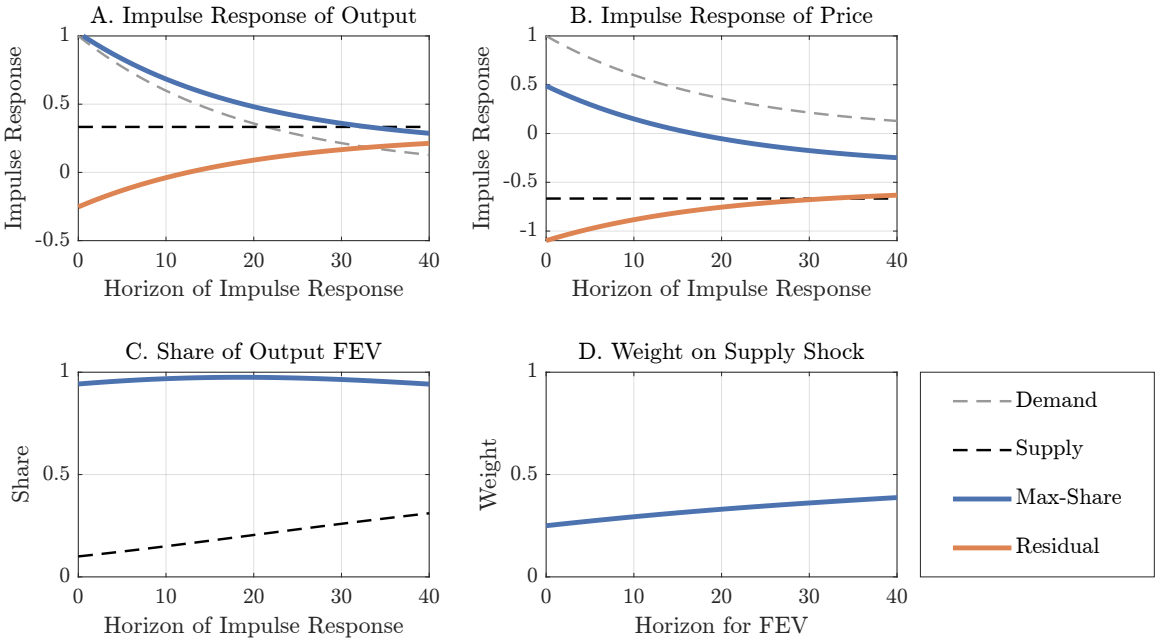


Figure 4: Permanent technology shock in supply and demand example identified via max-share in the time domain. **Panels A and B:** True and identified impulse responses; **Panel C:** Contribution to output FEV; **Panel D:** Weight that max-share identified technology shock places on true supply shock. Dashed lines indicate true responses; blue and orange solid lines correspond to identified max-share and residual shocks, respectively.

shock only places a weight of 0.39 on the supply shock, as seen in Panel D. While the weight is increasing in H and eventually converges to 1 as $H \to \infty$, the improvement is relatively slow, with the weight only increasing by 0.14 between $H = 0$ and $H = 40$. These results reflect the continued effect of the responses at short horizons.

We also plot impulse responses to what we label as the "residual" shock, which is the untargeted shock implied by max-share identification. Even though each of the true shocks produces a strictly positive response from output, the corresponding response to the residual shock is negative for short horizons and then becomes positive after horizon 12. This ensures orthogonality to the max-share shock, which is by construction. Even in models with more than two shocks, the orthogonality condition imposes restrictions on the untargeted responses that can potentially imply counterintuitive results.[20]

[20]In the Online Appendix, we identify a main business cycle shock using max-share at business cycle frequencies, following Angeletos et al. (2020), and highlight similar implications for the residual shock.

6.3 Medium-Scale New Keynesian Model

We also consider data generated from the medium-scale New Keynesian model of Smets and Wouters (2007), with parameters set to the posterior mode reported in the paper. We estimate a VAR with four lags using the variables used by Smets and Wouters (2007) for estimation—GDP growth, inflation, the interest rate, consumption growth, investment growth, the wage rate, and hours. To minimize estimation uncertainty, we use a long sample of 10^6 periods. We identify a main business cycle shock using max-share in the frequency domain, targeting GDP growth at frequency band $\Omega = \left[\frac{2\pi}{32}, \frac{2\pi}{6}\right]$, as in Angeletos et al. (2020).

In Figure 5, we plot the results together with the impulse responses and FEV decomposition from the data-generating process (DGP). The main business cycle shock has a small effect on inflation, consistent with empirical findings in Angeletos et al. (2020) and Section 7 below. However, none of the shocks in the DGP behave like the main business cycle shock. Panel B shows that almost all the true shocks produce larger (cumulative) inflation responses than the business cycle shock beyond the one year horizon; the small price response to the business cycle shock arises from the true responses to different shocks offsetting each other. Moreover, Panel C shows that the business cycle shock explains substantially more of GDP growth variation than any of the individual DGP shocks. Nevertheless, the small contribution to the FEV of inflation in Panel D does coincide with the price and wage markup shocks being the main drivers of inflation but contributing little to GDP growth in the model.

Even though we cannot recover the true shocks from the DGP since it is not invertible, we can obtain approximate weights on them using Theorem 2 and projecting the GDP growth impulse responses from the DGP on the max-share impulse response. The max-share shock does not pick out any single shock. Instead, the TFP, exogenous spending, investment-specific, and monetary shocks each account for about 1/5 or more of the weight. The even weights are consistent with Theorem 1. Panel A shows that the GDP impulse responses have similar shapes, thus violating orthogonality. In addition, Panel C shows that no single shock dominates the variation at business cycle frequencies, suggesting that we do not have the case of a large gap between eigenvalues described in the discussion of Propositions 1 and 2, making the deviation from orthogonality an important determinant of the max-share shock. The weights are also broadly in line with the relative contributions to the FEV at the target frequency band of $\left[\frac{2\pi}{32}, \frac{2\pi}{6}\right]$.

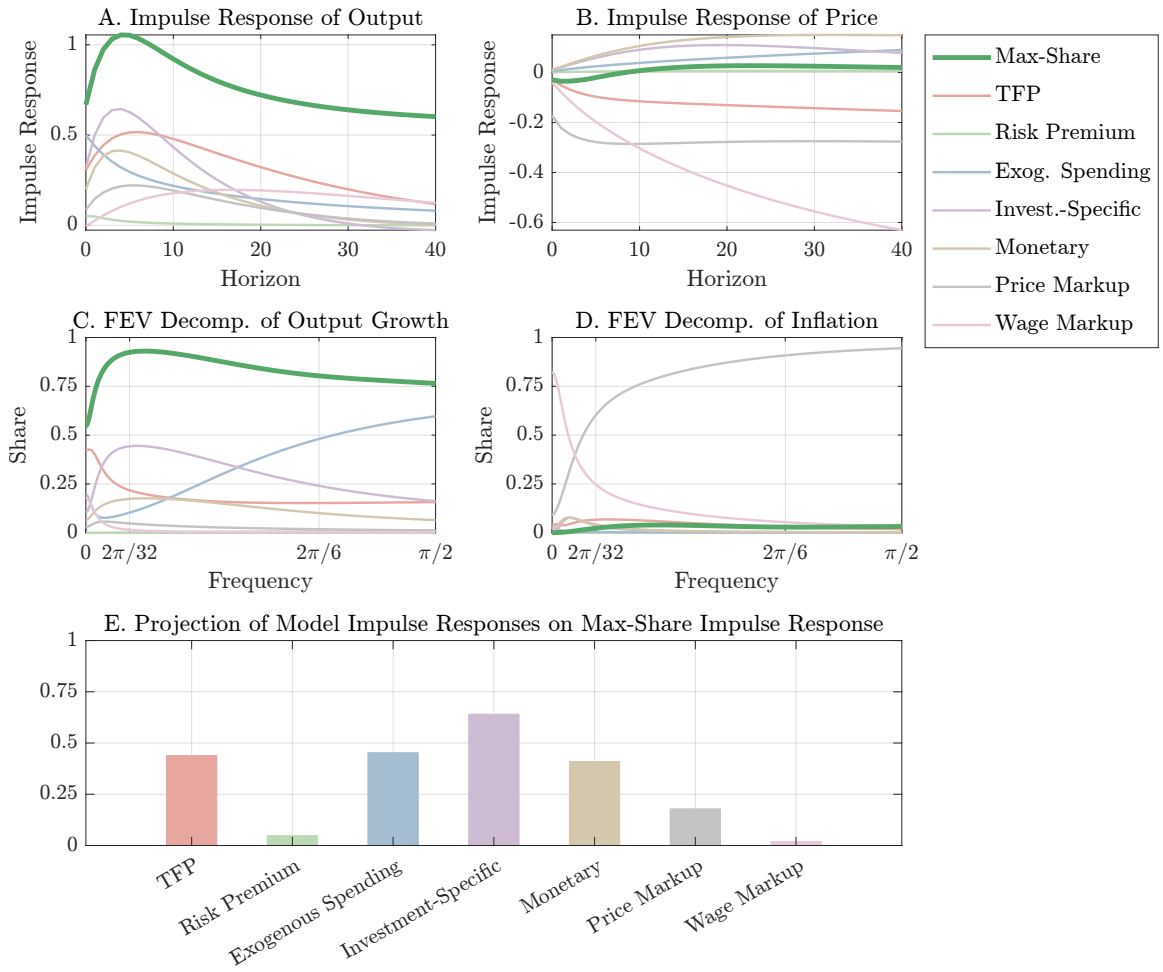


Figure 5: Results from Smets and Wouters (2007) simulation. **Panels A and B:** VAR and DGP impulse responses; **Panels C and D:** Contributions to FEV in VAR and DGP; **Panel E:** Weight that max-share identified main business cycle shock places on each true shock, as implied by projections. Dark green line corresponds to max-share shock; light colored lines correspond to shocks from DGP.

Our findings highlight the limitations of max-share identification. As emphasized by Giannone et al. (2019), the impulse responses to the max-share shock can summarize how various variables tend to respond to shocks that affect the target variable on average, but they do not generally capture a single structural shock. Instead of modeling a single shock to match the VAR's max-share shock, one can validate a structural model by identifying an analogous max-share shock in that model and checking if the max-share shocks in the VAR and structural model produce impulse responses with similar features. For example, Angeletos et al. (2020) argue that even

if the max-share shock consists of multiple true structural shocks, the empirical finding of "interchangeability," i.e., similar responses with different variables targeted by max-share, is something researchers should seek to replicate in DSGE models.

7 Empirical Applications

We now consider two well-known empirical applications of max-share identification: TFP news shocks following Kurmann and Sims (2021) and a main business cycle shock as in Angeletos et al. (2020). In addition, we identify a TFP surprise shock by using recursive identification with observed TFP ordered first in the VAR. We apply Theorem 2 to quantify how much the identified TFP max-share shock is contaminated by the TFP surprise shock and how much the identified main business cycle shock is contaminated by the TFP news max-share shock.

7.1 Data, Estimation, and Identification

We follow Kurmann and Sims (2021) and estimate an 8-variable VAR with utilization-adjusted TFP from Fernald (2014), real consumption per capita, real investment per capita, real GDP per capita, hours per capita, GDP deflator inflation, the federal funds rate, and the S&P 500 index. Angeletos et al. (2020) use a similar VAR but also include measures of labor productivity, the labor share, and unemployment while omitting the S&P 500 index. All variables except inflation and the federal funds rate are in log-levels. Our sample period is 1960Q1 through 2019Q4. We use a Minnesota prior with tightness parameter chosen to maximize the marginal likelihood.[21]

Building on a large literature (e.g., Beaudry and Portier, 2006; Barsky and Sims, 2011; Schmitt-Grohé and Uribe, 2012) that seeks to identify shocks that affect future productivity without being related to current or past fundamentals, we separately identify TFP news and surprise shocks. The standard assumption in the literature is that the TFP news shock affects TFP with a lag as it takes time for the information to diffuse. In contrast, TFP can respond to the surprise shock on impact. Since the arrival of news will have an impact on future TFP through, for instance, technology

[21] There are three main differences relative to Kurmann and Sims (2021). First, we use the 2023 vintage of the TFP series. Second, we have a longer sample period. Third, we use different hyperparameters for the Minnesota prior. None of these materially affect our main conclusions.

diffusion, Barsky and Sims (2011) and Kurmann and Sims (2021) propose using max-share that targets TFP at a relatively long horizon to identify the news shock. We follow Kurmann and Sims (2021) and use max-share in the time domain, targeting TFP with $\mathcal{H} = \{0, \ldots, 40\}$. To identify a TFP surprise shock, we use recursive identification with TFP ordered first. This is consistent with an approach of interpreting our measure of TFP as an instrumental variable and identifies the surprise shock as being the one that explains all residual variation in TFP after controlling for lagged variables.

As in Section 6.3, identification of the business cycle shock is based on work by Angeletos et al. (2020). We follow their approach and use max-share in the frequency domain, targeting GDP at frequency band $\Omega = \left[\frac{2\pi}{32}, \frac{2\pi}{6}\right]$.[22] Angeletos et al. (2020) argue that this "main business cycle shock" resembles a demand shock that has minimal effect on inflation. Since TFP has a relatively small response to their business cycle shock, they rule out the role of supply shocks. They then argue that these responses should inform business cycle models with single shocks. While recent papers have questioned these conclusions (Bianchi et al., 2023; Forni et al., 2024; Granese, 2024), they do so through alternative statistical models. In contrast, we maintain the VAR structure used by Angeletos et al. (2020) and focus on the use of max-share.

Since the three shocks are identified independently, it is possible that the true structural shock(s) picked up by one identification scheme are also embedded in another of the identified shocks. However, economic theory suggests that they should be distinct. TFP news and surprise shocks are typically assumed to be independent and distinguishable by the response of TFP on impact. The business cycle shock, if interpreted as a pure demand shock, should not be related to TFP shocks. Theorem 2 gives us a way to test these hypotheses.

7.2 Results

Figure 6 shows the impulse responses to the three identified shocks. Despite the differences in the details of the reduced form VAR, the TFP news and business cycle shocks closely resemble those in Kurmann and Sims (2021) and Angeletos et al. (2020), respectively. However, with the identification conditions of Theorem 1 in mind, visual inspection of these responses already suggests that the two max-share

[22] Angeletos et al. (2020) use the same frequency band but target unemployment as their benchmark. However, they show that both target variables give similar results.

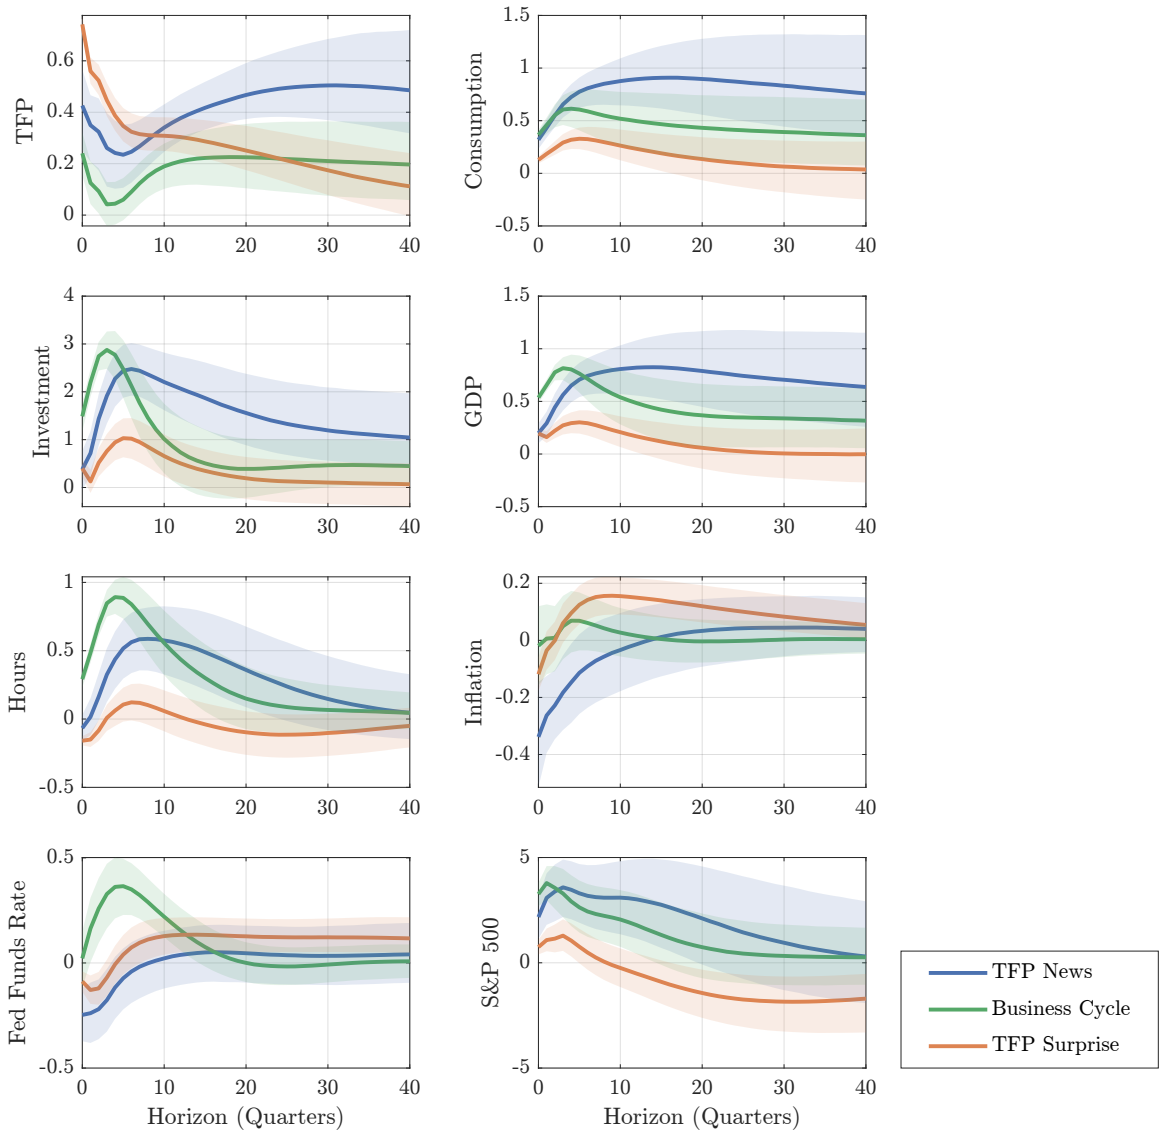


Figure 6: Posterior estimates for impulse responses to identified shocks. **Solid lines:** Median response; **Shaded regions:** 68% error bands.

shocks are likely contaminated.[23]

As we argued previously, the fact that the TFP news and surprise shocks are both strictly positive over the horizons of interest implies that orthogonality is not satisfied. The violations are made more severe by the persistence of the TFP surprise shock—the ratio of the TFP response at horizon 40 to horizon 0 is similar to that of an AR(1) with persistence 0.95. Moreover, the response to the TFP news shock has

[23]The Online Appendix provides further evidence from the FEV decomposition of the variables.

a hockey stick shape, with a substantial response at short horizons.[24]

While the analysis for the business cycle shock is less accessible because it is identified in the frequency domain, we observe that it produces a highly persistent response in GDP that is mirrored by the TFP news shock. While the shape of the responses over the short and medium run differ, the persistence likely plays an important role for the phase difference between the responses. In particular, the transfer function, Γ, in (8) is a sum over all horizons $\mathbb{Z}_{\geq 0}$. Since $|e^{-i\omega h}| = 1$, longer horizons will be a key determinant of Γ, thus playing an important role in the orthogonality violations through the phase difference between the responses.

We use Theorem 2 to quantify the contamination more formally, defining $\hat{\beta} \equiv \frac{\langle \psi^*, \hat{\psi} \rangle}{\langle \psi^*, \psi^* \rangle}$ and taking

$$\mathcal{C} \equiv \frac{\hat{\beta}}{\sqrt{1 - \hat{\beta}^2}} \tag{40}$$

as a measure of contamination. If $\hat{\psi}$ arises from a single structural shock $\hat{j}$, then this ratio corresponds to $\theta_{\hat{j}} / \sqrt{1 - \theta_{\hat{j}}^2}$, where the denominator constitutes the upper bound on $|\theta_1|$ (taking Shock 1 to be the target of max-share as before). Defining

$$\zeta \equiv \frac{\hat{\beta}}{\hat{\beta} + \sqrt{1 - \hat{\beta}^2}}$$

to be the fraction of the weight max-share places on the (combination of) shock(s) generating $\hat{\psi}$, we have that $\mathcal{C} = \frac{\zeta}{1-\zeta}$. In particular, if max-share does isolate the true structural shock of interest ($\zeta = 0$), then $\mathcal{C} = 0$. If half the max-share weight is on the (combination of) shock(s) generating $\hat{\psi}$ ($\zeta = 0.5$), then $\mathcal{C} = 1$.

Figure 7 plots the posterior for (40). With both the TFP news and business cycle shocks, max-share indeed places non-trivial weight on shocks that should be distinct. For the TFP news shock, the posterior for $\mathcal{C}$ peaks around 0.6; for the business cycle shock, the posterior peaks around 0.4. As measured by the ζ, over a third and a quarter, respectively, of the TFP news and business cycle max-share shocks therefore consist of shocks that a valid identification scheme would have excluded.

The discussion above centered only on the responses of the target variables, TFP

[24] Kurmann and Sims (2021) find a similarly shaped impulse response, but with a smaller initial effect. One could avoid this feature by imposing a zero response on impact as done, for instance, by Barsky and Sims (2011) and Görtz et al. (2022).

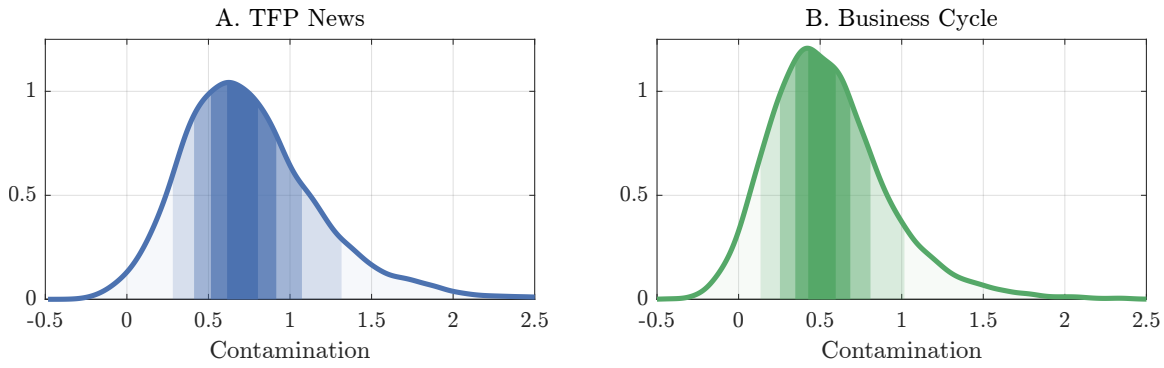


Figure 7: Posterior distribution for contamination of TFP news (left) and business cycle (right) max-share shocks by identified TFP surprise and TFP news shocks, respectively. Shaded regions indicate deciles.

and GDP. This is motivated by our theoretical results and the primitive problems, (4) and (6), neither of which reference the untargeted variables. In Figure 6, there are untargeted variables for which the responses to TFP news and business cycle shocks differ substantially. For example, while the TFP news shock leads to a decline in inflation (consistent with supply shocks in theory), the business cycle shock leads to a small but statistically insignificant increase in inflation. Similarly, the two shocks produce opposite responses in the federal funds rate. Nevertheless, we find formal evidence that the business cycle shock contains the shock(s) picked up by the TFP news shock identification. Therefore, the small contributions to TFP and inflation by the main business cycle shock do not imply that we should rule out a role for shocks that affect TFP or inflation in driving the business cycle. More generally, the interpretation of the untargeted impulse responses to the max-share shock depends on whether the conditions in Theorem 1 are satisfied.

While our results urge caution in interpreting the max-share identified TFP news and main business cycle shocks, they do not rule out their usefulness. For example, if one is willing to assume that the measurement error emphasized by Kurmann and Sims (2021) has a relatively small and transitory effect on measured TFP, then the max-share TFP news shock plausibly captures the overall combination of TFP news and surprise shocks (and any other shocks) that are most important for driving TFP at a medium horizon. This interpretation can inform how we use the max-share impulse responses to model TFP. Specifically, the model's responses to the TFP news shock should not directly match the max-share impulse responses. Instead, the max-share impulse responses and contamination measure can inform how persistent and

large the TFP surprise shock is relative to the news shock. This echoes our discussion of the main business cycle shock in Section 6.3.

8 Conclusion

For max-share to correctly identify a target shock, we show that the following necessary and sufficient conditions are needed: (i) the impulse response of the target variable to the shock of interest is orthogonal to the corresponding responses to all other shocks and (ii) the response to the target shock is large relative to combinations of responses to the untargeted shocks, as summarized by the corresponding eigenvalues. These conditions are hard to satisfy in practice, and violations can lead max-share to severely misidentify the shock of interest. With knowledge of the responses of the target variable to untargeted shocks, we have a straightforward way to measure the degree of contamination.

The stringency of the conditions does not preclude the usefulness of max-share for disciplining DSGE models. However, researchers should ensure that their DSGE model is congruent with the identification assumptions of the structural VAR. This requirement is common practice in the impulse response matching literature (e.g., Christiano et al., 2005; Fisher, 2006; Altig et al., 2011), where assumptions are made so that the DSGE models satisfy the VAR shock identification restrictions. In other words, a VAR estimated on data simulated from the DSGE and identified using the same restrictions should recover the shocks of interest. To compare the max-share identified shock to a single shock in a DSGE model, we need to argue that the orthogonality and relative size conditions survive the introduction of other shocks. If we are unable to make such a claim, we can still use the max-share impulse responses to validate or quantify the model through indirect inference, by showing that data generated from the DSGE generates a max-share shock with similar properties to the max-share shock in the VAR.

References

Altig, D., L. J. Christiano, M. Eichenbaum, and J. Linde (2011). Firm-Specific Capital, Nominal Rigidities and the Business Cycle. *Review of Economic Dynamics 14*(2), 225–247.

An, S. and F. Schorfheide (2007). Bayesian Analysis of DSGE Models. *Econometric Reviews 26*(2-4), 113–172.

Angeletos, G.-M., F. Collard, and H. Dellas (2020). Business-Cycle Anatomy. *American Economic Review 110*(10), 3030–3070.

Arias, J. E., J. F. Rubio-Ramírez, and D. F. Waggoner (2018). Inference Based on Structural Vector Autoregressions Identified With Sign and Zero Restrictions: Theory and Applications. *Econometrica 86*(2), 685–720.

Barsky, R. B. and E. R. Sims (2011). News Shocks and Business Cycles. *Journal of Monetary Economics 58*(3), 273–289.

Basu, S., G. Candian, R. Chahrour, and R. Valchev (2025). Risky Business Cycles. Working paper.

Bauer, M. D. and E. T. Swanson (2023a). A Reassessment of Monetary Policy Surprises and High-Frequency Identification. *NBER Macroeconomics Annual 37*(1), 87–155.

Bauer, M. D. and E. T. Swanson (2023b, March). An Alternative Explanation for the "Fed Information Effect". *American Economic Review 113*(3), 664–700.

Beaudry, P. and F. Portier (2006). Stock Prices, News, and Economic Fluctuations. *American Economic Review 96*(4), 1293–1307.

Ben Zeev, N. and H. Khan (2015). Investment-Specific News Shocks and US Business Cycles. *Journal of Money, Credit and Banking 47*(7), 1443–1464.

Bianchi, F., G. Nicolò, and D. Song (2023). Inflation and Real Activity Over the Business Cycle. Working Paper 31075, National Bureau of Economic Research.

Blanchard, O. J. and D. Quah (1989). The Dynamic Effects of Aggregate Demand and Supply Disturbances. *American Economic Review 79*(4), 655–673.

Brennan, C. M., M. M. Jacobson, C. Matthes, and T. B. Walker (2024). Monetary Policy Shocks: Data or Methods? Working paper.

Caldara, D., C. Fuentes-Albero, S. Gilchrist, and E. Zakrajšek (2016). The Macroeconomic Impact of Financial and Uncertainty Shocks. *European Economic Review 88*, 185–207.

Carriero, A. and A. Volpicella (2024). Max Share Identification of Multiple Shocks: An Application to Uncertainty and Financial Conditions. *Journal of Business & Economic Statistics*, 1–13.

Cascaldi-Garcia, D. and A. B. Galvao (2021). News and Uncertainty Shocks. *Journal of Money, Credit and Banking 53*(4), 779–811.

Chahrour, R., V. Cormun, P. De Leo, P. Guerrón-Quintana, and R. Valchev (2024). Exchange Rate Disconnect Revisited. Working paper.

Chen, Y. and D. Liu (2019). How Does Government Spending News Affect Interest Rates? Evidence from the United States. *Journal of Economic Dynamics and Control 108*, 103747.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy 113*(1), 1–45.

Del Negro, M. and F. Schorfheide (2004). Priors from General Equilibrium Models for VARs. *International Economic Review 45*(2), 643–673.

DiCecio, R. and M. Owyang (2010). Identifying Technology shocks in the Frequency Domain. Federal Reserve Bank of St. Louis Working Paper 2010-025A.

Dieppe, A., N. Francis, and G. Kindberg-Hanlon (2021). The Identification of Dominant Macroeconomic Drivers: Coping with Confounding Shocks. ECB Working Paper 2534.

Engle, R. F., D. F. Hendry, and J.-F. Richard (1983). Exogeneity. *Econometrica*, 277–304.

Faust, J. (1998). The Robustness of Identified VAR Conclusions About Money. In *Carnegie-Rochester Conference Series on Public Policy*, Volume 49, pp. 207–244.

Fernald, J. (2014). A Quarterly, Utilization-Adjusted Series on Total Factor Productivity. Federal Reserve Bank of San Francisco Working Paper 2012-19.

Fernández-Villaverde, J., J. F. Rubio-Ramírez, T. J. Sargent, and M. W. Watson (2007). ABCs (and Ds) of Understanding VARs. *American Economic Review 97*(3), 1021–1026.

Fève, P. and A. Guay (2019). Sentiments in SVARs. *The Economic Journal 129*(618), 877–896.

Fisher, J. D. (2006). The Dynamic Effects of Neutral and Investment-Specific Technology Shocks. *Journal of Political Economy 114*(3), 413–451.

Forni, M., L. Gambetti, A. Granese, L. Sala, S. Soccorsi, et al. (2024). An American Macroeconomic Picture: Supply and Demand Shocks in the Frequency Domain. *American Economic Journal: Macroeconomics*, 1–67.

Francis, N. and G. Kindberg-Hanlon (2022). Signing Out Confounding Shocks in Variance-Maximizing Identification Methods. In *AEA Papers and Proceedings*, Volume 112, pp. 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 Economics and Statistics 96*(4), 638–647.

Francis, N. and V. A. Ramey (2009). Measures of per Capita Hours and Their Implications For the Technology-Hours Debate. *Journal of Money, Credit and Banking 41*(6), 1071–1097.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle: an Introduction to the New Keynesian Framework and Its Applications*. Princeton University Press.

Gertler, M. and P. Karadi (2015). Monetary Policy Surprises, Credit Costs, and Economic Activity. *American Economic Journal: Macroeconomics 7*(1), 44–76.

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 61*, 137–173.

Görtz, C., C. Gunn, and T. A. Lubik (2022). Is There News in Inventories? *Journal of Monetary Economics 126*, 87–104.

Görtz, C., J. D. Tsoukalas, and F. Zanetti (2022). News Shocks Under Financial Frictions. *American Economic Journal: Macroeconomics 14*(4), 210–243.

Gospodinov, N., A. María Herrera, and E. Pesavento (2013). Unit Roots, Cointegration, and Pretesting in VAR Models. In *VAR Models in Macroeconomics–New Developments and Applications: Essays in Honor of Christopher A. Sims*, pp. 81–115. Emerald Group Publishing Limited.

Granese, A. (2024). Two Main Business Cycle Shocks are Better than One. Working paper.

Guay, A., F. Pelgrin, and S. Surprenant (2024). Max Share Identification for Structural VARs in Levels: There is No Free Lunch! Working paper.

Ingram, B. F. and C. H. Whiteman (1994). Supplanting the 'Minnesota' Prior: Forecasting Macroeconomic Time Series Using Real Business Cycle Model Priors. *Journal of Monetary Economics 34*(3), 497–510.

Kilian, L., M. Plante, and A. W. Richter (2023). Jointly Estimating Macroeconomic News and Surprise Shocks. Working paper.

King, R. G., C. I. Plosser, and S. T. Rebelo (1988). Production, Growth and Business Cycles: I. The Basic Neoclassical Model. *Journal of Monetary Economics 21*(2-3), 195–232.

Kurmann, A. and C. Otrok (2013). News Shocks and the Slope of the Term Structure of Interest Rates. *American Economic Review 103*(6), 2612–2632.

Kurmann, A. and E. Sims (2021). Revisions in Utilization-Adjusted TFP and Robust Identification of News Shocks. *Review of Economics and Statistics 103*(2), 216–235.

Levchenko, A. A. and N. Pandalai-Nayar (2020). TFP, News, and "Sentiments": The International Transmission of Business Cycles. *Journal of the European Economic Association 18*(1), 302–341.

McKay, A. and C. K. Wolf (2023). What Can Time-Series Regressions Tell Us About Policy Counterfactuals? *Econometrica 91*(5), 1695–1725.

Meier, M. and T. Reinelt (2024). Monetary Policy, Markup Dispersion, and Aggregate TFP. *Review of Economics and Statistics 106*(4), 1012–1027.

Mertens, K. and M. O. Ravn (2013). The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States. *American Economic Review 103*(4), 1212–1247.

Miranda-Agrippino, S. and G. Ricco (2021). The Transmission of Monetary Policy Shocks. *American Economic Journal: Macroeconomics 13*(3), 74—107.

Miyamoto, W., T. L. Nguyen, and H. Oh (2023). In Search of Dominant Drivers of the Real Exchange Rate. *Review of Economics and Statistics*, 1–50.

Moench, E. and S. Soofi-Siavash (2022). What Moves Treasury Yields? *Journal of Financial Economics 146*(3), 1016–1043.

Mumtaz, H., G. Pinter, and K. Theodoridis (2018). What do VARs Tell Us About the Impact of a Credit Supply Shock? *International Economic Review 59*(2), 625–646.

Noh, E. (2018). Impulse-Response Analysis with Proxy Variables. Working paper, University of California San Diego.

Phillips, P. C. (1998). Impulse Response and Forecast Error Variance Asymptotics in Nonstationary VARs. *Journal of Econometrics 83*(1-2), 21–56.

Plagborg-Møller, M. (2019). Bayesian Inference on Structural Impulse Response Functions. *Quantitative Economics 10*(1), 145–184.

Plagborg-Møller, M. and C. K. Wolf (2021). Local Projections and VARs Estimate the Same Impulse Responses. *Econometrica 89*(2), 955–980.

Priestley, M. B. (1981). *Spectral Analysis and Time Series*. New York: Academic Press.

Sargan, J. D. (1958). The Estimation of Economic Relationships Using Instrumental Variables. *Econometrica*, 393–415.

Schmitt-Grohé, S. and M. Uribe (2012). What's News in Business Cycles. *Econometrica 80*(6), 2733–2764.

Sims, C. A. (1980). Macroeconomics and Reality. *Econometrica*, 1–48.

Smets, F. and R. Wouters (2007). Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review 97*(3), 586–606.

Stock, J. H. and M. W. Watson (1999). Business Cycle Fluctuations in US Macroeconomic Time Series. In *Handbook of Macroeconomics*, Volume 1, Chapter 1, pp. 3–64. Elsevier.

Stock, J. H. and M. W. Watson (2016). Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics. In *Handbook of Macroeconomics*, Volume 2, pp. 415–525. Elsevier.

Stock, J. H. and M. W. Watson (2018). Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments. *The Economic Journal 128*(610), 917–948.

Uhlig, H. (2004a). Do Technology Shocks Lead to a Fall in Total Hours Worked? *Journal of the European Economic Association 2*(2-3), 361–371.

Uhlig, H. (2004b). What Moves GNP? Working paper.

Uhlig, H. (2005). What Are the Effects of Monetary Policy on Output? Results From an Agnostic Identification Procedure. *Journal of Monetary Economics 52*(2), 381–419.

Watson, M. (2001). Time Series: Cycles. In N. J. Smelser and P. B. Baltes (Eds.), *International Encyclopedia of the Social & Behavioral Sciences*, pp. 15714–15721. Oxford: Pergamon.