Working Paper Series

Max-Share Misidentification

WP 25-02

Liyu Dou
Singapore Management University
CUHK-Shenzhen

Paul Ho
Federal Reserve Bank of Richmond

Thomas A. Lubik
Federal Reserve Bank of Richmond

FEDERAL RESERVE BANK OF RICHMOND®
Richmond ▪ Baltimore ▪ Charlotte

Max-Share Misidentification*

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

[1]Singapore Management University and CUHK-Shenzhen

[2,3]Federal Reserve Bank of Richmond

May 23, 2025

Abstract

Valid max-share identification requires necessary and sufficient conditions that are hard to satisfy 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. We theoretically characterize consequences of local and global violations to these conditions. 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. Empirically, the TFP news and 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, Felipe Schwartzman, and participants at seminars and conferences 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 to 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 in that it not only yields point identification unlike sign restrictions, but it appears to rest 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, in contrast to potentially controversial zero restrictions for recursive identification or exogeneity requirements for instrumental variables.

We show that max-share identification is problematic. We formally derive the necessary and sufficient conditions for max-share identification to be valid and demonstrate that these are substantially more stringent than previously perceived in the literature. The 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 as capturing the covariance across impulse responses of the target variable to the true structural shocks. Instead of this covariance being of random variables over observations, it 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 condition is *orthogonality* of the impulse responses, which corresponds to Ξ being a block diagonal matrix with the off-diagonal terms in the first row and column being zero. This 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 $h-1$. 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. While the interpretation is less transparent in the frequency domain, we show that orthogonality is violated if the responses of the target shock and some untargeted shock resemble the impulse responses of stationary AR(1) models. Since orthogonality has to be satisfied for the full set of shocks, it requires the econometrician to have 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 needs to 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 not only be large relative to the response to each of the other individual shocks, but to 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 will be large, potentially causing 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 $h \geq 0$.

We illustrate these conditions by means of four stylized examples in Figure 1. In each case, the econometrician seeks to identify shock 1 and knows that this shock produces the largest response at horizon 1. A common practice in the literature would then be to use this knowledge as justification for using max-share with target horizon 1 to identify the shock. The top left panel presents the ideal case—max-share perfectly 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 of shock 2 only peaks at 0.9, the relative size condition is also satisfied.

The subsequent panels show how slight deviations threaten identification. On the top right, orthogonality is violated as we now increase the horizon 1 response of 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. The bottom left panel instead 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, so that the order of the

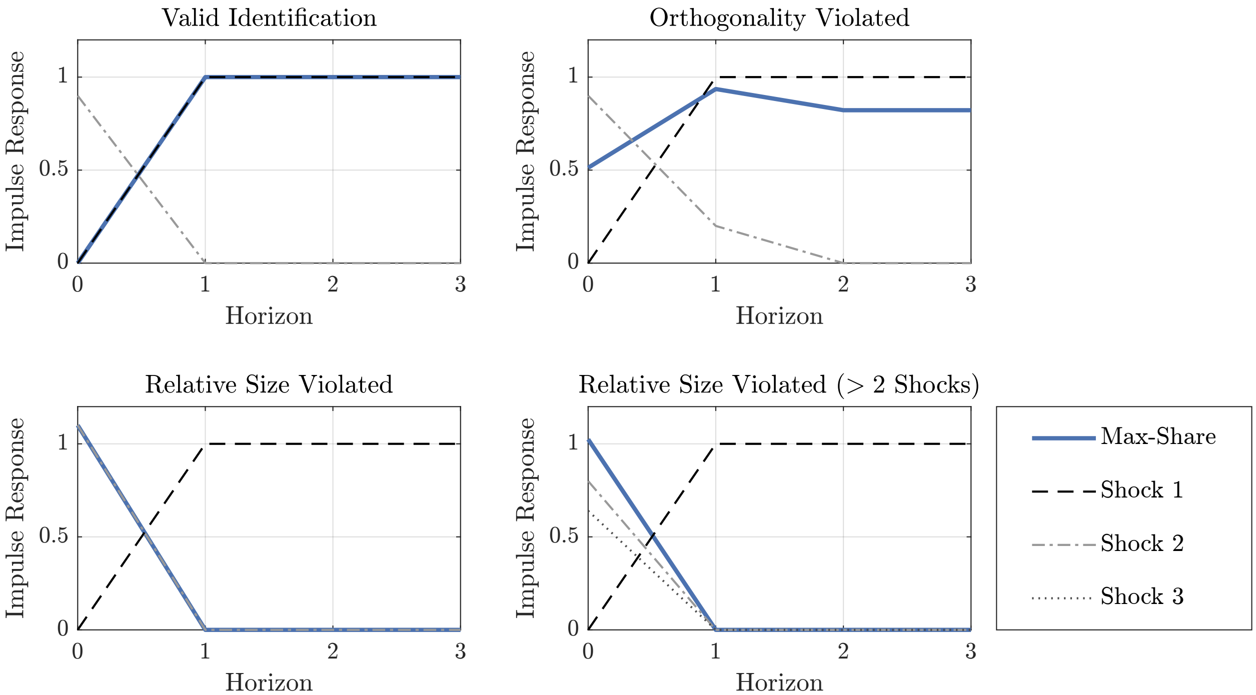

Figure 1: Impulse responses to max-share identified and true shocks in four stylized examples. Solid blue line corresponds to max-share shock; remaining lines correspond to true shocks.

eigenvalues is now flipped. 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, the bottom right panel 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 present various examples to show that these issues are pervasive across common impulse responses in the literature.

More generally, 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.

Finally, we show that one can obtain a lower bound on how much the max-share identified shock is contaminated in practice. Suppose that besides the max-share shock, one also observes an untargeted shock. The weight that max-share places on the untargeted shock is the projection of the impulse response of the target variable to the max-share shock on the corresponding response to the observed shock. Consequently, we can place an upper bound on the max-share weight on the true targeted shock.

In light of our results, we revisit the shocks identified by Kurmann and Sims (2021) and Angeletos et al. (2020). Using 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.

Related Literature. The max-share problem was initially introduced in the time domain in seminal work by Faust (1998), whose original goal was 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 pioneering contributions, the max-share approach has been used to identify a vast array of shocks, including 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), risk premium shocks (Basu, Candian, Chahrour, and Valchev, 2025), and exchange rate shocks (Miyamoto, Nguyen, and Oh, 2023; Chahrour, Cormun, De Leo, Guerrón-Quintana, and Valchev, 2024). Related work by Barsky and Sims (2011), Kurmann and Otrok (2013), and Ben Zeev and Khan (2015) 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 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 narrowly 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. In response, Francis and Kindberg-Hanlon (2022) use sign restrictions to address this issue. Similarly, Kilian, Plante, and Richter (2023) point out that the use of max-share and related approaches by Barsky and Sims (2011), Kurmann and Sims (2021), and Dieppe et al. (2021) to identify news shocks may produce misleading results and propose the use of instrumental variables 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 being 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 argue that in the presence of unit or near-unit roots reduced form impulse responses are inconsistent, threatening the validity of max-share. Throughout our analysis, we instead assume 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 presents the optimization problems for max-share identification. Section 3 presents our main theoretical results in a unified framework for the time and frequency domains. Section 4 specializes these insights for the time and frequency domains and draws comparison with other modes of VAR identification. Sections 5 and 6 illustrate our insights using numerical examples and an empirical application, respectively. Section 7 concludes.

2 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 $E[\varepsilon_t \varepsilon_t'] = I$. We can write 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.[1] 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.

2.1 Time Domain

The most common approach to max-share identification works in the time domain. To identify Ψ_h, it considers the problem:

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

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

The typical use of (3) in max-share identification, following Faust (1998), corresponds to $\mathcal{H} = \{0, 1, \ldots, H\}$. This maximizes the contribution of the target shock to the FEV of variable i. Specifically, $\widetilde{\Psi}_h \theta \theta' \widetilde{\Psi}_h'$ in the numerator is an $N \times N$ matrix that 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 alone. This approach is the 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). The single horizon problem with $\mathcal{H} = \{0\}$ 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:*

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

[1] We have $\Psi_h = (\mathbb{B}^h)_{1:N,1:N} C$, where $\mathbb{B}$ is the autoregressive coefficient in the companion form of (1) and $(\mathbb{B}^h)_{1:N,1:N}$ denotes the upper-left $N \times N$ submatrix of $\mathbb{B}^h$.

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 $\theta'\theta = 1$.

2.2 Frequency Domain

An alternative approach proposed by DiCecio and Owyang (2010) and Angeletos et al. (2020) 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 and not be contaminated by overly high or low frequency fluctuations in the data.

The frequency domain max-share approach solves the following:

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

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

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

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

$\Omega \subseteq [-\pi, \pi]$ 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 MA representation (2) with the structural shocks ε_{t-h} replaced by the Euler coefficient $e^{-i\omega h}$, 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 $\mathrm{Re}\left[\int_{\omega\in\Omega} \Upsilon(\omega)\, d\omega\right]$ as shown in Lemma 1.[3]

3 Max-Share Misidentification: General Results

For exposition, we henceforth take $\widetilde{\Psi}_h = \Psi_h$, so that we can interpret the jth entry of the solution to (3) as the weight that the max-share shock places on the jth true shock.[4]

[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. In what follows, we maintain the high-level assumption that $\sum_{h=0}^{\infty} |\Psi_{h,ij}| < \infty$ for $i, j = 1, \ldots, N$. 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.

3.1 Unifying the Time and Frequency Domain Problems

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

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

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. Intuitively, it can be thought of as a generalization of a covariance matrix in the case of a vector of random variables. We explain in the Online Appendix that it is rare for Ξ to have repeated eigenvalues. In particular, this implies that (9) will generally have a unique solution.

The matrix Ξ will depend on the reduced form parameters $\{\{B_\ell\}_\ell^L, \Sigma\}$ as well as $\mathcal{H}$ in the time domain and Ω in the frequency domain. The h-th element of the vector ψ_j contains the impulse response of the target variable to shock j at the h-th horizon in some set of horizons $\mathfrak{H}$. The solution for θ is the principal eigenvector (i.e., the eigenvector associated with the largest eigenvalue) of Ξ.

Suppressing the dependence on the reduced form parameters 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:

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

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

from (6) in the frequency domain, where $\Gamma_{i,j}^{\mathrm{Re}}(\omega)$ and $\Gamma_{i,j}^{\mathrm{Im}}(\omega)$ are the real and imaginary parts of the (i,j) element of the transfer function, $\Gamma(\omega)$.[5] 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, but of impulse responses and not variables. In other words, even though the structural shocks might be orthogonal, the off-diagonal elements of Ξ are large (relative to the diagonals) if the shocks have similar dynamic effects on the target variable i. In practice, column (or row) j corresponds to the response of the

[5] 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 $\{H_k\}_{k=1}^K$.

target variable to the true shock j. We provide further intuition when we discuss the structure of Ξ^{time} and Ξ^{freq} in detail in Section 4. Henceforth, we assume without loss of generality that the true target structural shock is ordered first.

Valid identification then corresponds to (9) having solution $\theta = \delta_1$, so that all the weight is placed on the true target structural shock. Since the solution is the principal eigenvector of Ξ, we now study conditions that a generic Gram matrix Ξ must satisfy to yield a principal eigenvector of δ_1. In addition, we describe the behavior of the principal eigenvector as we deviate from these conditions. Finally, we show that an econometrician observing (or externally identifying) an untargeted shock can bound the degree of contamination. All proofs are in the Online Appendix.

3.2 Necessary and Sufficient Conditions for the Validity of Max-Share

Denote the *direct sum* of two square matrices A and B by $A \oplus B$, i.e.,

$$A \oplus B \equiv \begin{bmatrix} A & 0 \\ 0 & B \end{bmatrix}.$$

In addition, let $\lambda_{\max}(B)$ denote the largest eigenvalue of B 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) $\Xi_{1,1} > \lambda_{\max}(\Xi_{2:N,2:N})$, i.e., the (first) spectral gap of Ξ is strictly positive.*

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 exact economic 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 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. The orthogonality condition 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, the top right panel 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 many empirical applications of max-share.

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 bottom right panel 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.

Intuitively, for max-share identification to be justified it is not enough that the target impulse response is larger than the response of the target variable to each of the remaining shocks individually. Instead, the target responses cannot be similar in shape to the responses to the untargeted shocks. For instance, if we attempted to identify TFP shocks by targeting output at a long but finite horizon, the identification conditions are likely to be violated since the response of output to TFP shocks shares some similarity in 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 the other shocks are fully dominated by the target shock over the target horizons or frequencies. This is the case for long-run identification but is unlikely in other settings.

Both conditions emphasize the neglected role of impulse response shapes in the existing max-share literature. We also note that the shape restrictions apply to the responses of the target variable to the full set of shocks jointly. In the context of the VAR (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.[6]

It is also common to claim that the target variable is driven by a single shock or to propose 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. However, 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)). As long as the identification restrictions are violated, max-share will overstate the importance of the targeted shock. Moreover, there is no guarantee that the response of untargeted variables to the max-share shock correspond

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

to any single true structural shock.

3.3 Perturbations and Global Bounds

To illustrate the behavior of the max-share approach when the identification conditions in Theorem 1 are not satisfied, we now provide local approximations and bounds for the solution to (9), focusing on deviations from the orthogonality condition.

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

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

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

$$\left[1, \sum_{j=2}^{N} \frac{w_{0j}'\nu}{\Xi_{1,1} - \lambda_{0j}} w_{0j}'\right]' + O(\|\nu\|^2), \tag{15}$$

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, $\partial\Xi$, is zero on the diagonal blocks and characterized by the vector ν off the diagonal blocks. Intuitively, we can think of this 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, (15), 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, then small deviations from orthogonality will have a relatively minor effect on identification.

While the solution to (9) is not generally analytical, we can derive bounds that apply globally.

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

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

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

$$\sin\alpha(\delta_1, \theta) \le \frac{2\,\|\nu\|}{\Xi_{1,1} - \lambda_{\max}(\Xi_{2:N,2:N})} \quad \textit{and} \quad \|\theta - \delta_1\| \le \frac{2^{3/2}\,\|\nu\|}{\Xi_{1,1} - \lambda_{\max}(\Xi_{2:N,2:N})}, \tag{16}$$

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 $\partial\Xi$ as Proposition 1, but no longer requires that $\|\nu\|$ is small. The two expressions in (16) 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 between the two largest eigenvalues, $\Xi_{1,1} - \lambda_{\max}(\Xi_{2:N,2:N})$, of Ξ. If the spectral gap is sufficiently large, then even nontrivial deviations from orthogonality will result in small degrees of contamination.[7]

3.4 Bounds on Contamination in Practice

We now 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. Specifically, if we separately identify an untargeted shock (and the associated impulse response), we can check if the identification conditions are satisfied. Moreover, we can obtain a measure of how much the max-share identified shock is contaminated by the externally identified shock. More generally, the untargeted shock can be identified using any identification scheme as long as it is plausible that it is 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 ends up capturing a combination of true untargeted shocks. The theorem below thus 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.

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*

[7] Consider a Blanchard and Quah (1989) economy where only one shock has a permanent effect on the variable of interest. With sufficiently long horizon, one could still 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 $\psi_{2,h} = 0.9^h$. Then for $\mathcal{H} = \{0, \ldots, 100\}$, we have $\|\psi_{1,\mathcal{H}}\|^2 / \|\psi_{2,\mathcal{H}}\|^2 = 16$ and max-share places 95% of the weight on Shock 1 even though orthogonality is violated with $\psi_{1,\mathcal{H}} \cdot \psi_{2,\mathcal{H}} / \|\psi_{1,\mathcal{H}}\|\,\|\psi_{2,\mathcal{H}}\| = 0.25$. Nonetheless, other drawbacks to long-run identification arising from low frequency fluctuations have been documented (Francis and Ramey, 2009; Gospodinov et al., 2013).

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:*

$$\langle \psi^*, \hat{\psi} \rangle = \langle \psi^*, \psi^* \rangle \sum_{j=2}^{N} \alpha_j \theta_j = \lambda_{\max}(\Xi) \sum_{j=2}^{N} \alpha_j \theta_j. \tag{17}$$

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

$$\theta_1^2 \leq 1 - \left(\frac{\langle \psi^*, \hat{\psi} \rangle}{\langle \psi^*, \psi^* \rangle} \right)^2 . \tag{18}$$

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.[8] Suppose we observe Shock 2 (or its impulse response), i.e., $\alpha_2 = 1$ and $\alpha_j = 0$ for $j > 2$. Then (17) 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 variable to the max-share shock, ψ^*. The bound (18) then follows, giving an upper bound on the weight that max-share identification is placing on the target shock. In a model with $N = 2$ shocks, the bound holds with equality.

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.[9]

As an alternative approach, one could first control for the observed shock before implementing max-share (Cascaldi-Garcia and Galvao, 2021; Basu et al., 2025). Imagine a setting where multiple untargeted shocks are observed. A conceptual challenge is that the number of variables, N, limits the number of shocks in the VAR. Each structural shock also likely stands in for a more granular

[8] 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.

[9] 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.

set of shocks.[10] 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). Therefore, including multiple shocks in implementing Theorem 2 or controlling for these shocks potentially precludes other underlying shocks. Theorem 2 should thus be viewed as an identification diagnostic rather than a way to fully deal with any contamination of the max-share shock.

3.5 Two Special Cases

We now present two special cases for which we can explicitly solve for the solution to (9).

Rank One. Suppose Ξ is rank one. This occurs when the impulse responses of the target variable to all shocks have same shape. It represents the most severe violation of orthogonality. In this case, the solution to (9) satisfies:

$$\theta_j \propto \sqrt{\Xi_{j,j}}. \tag{19}$$

In other words, the contamination to the identified shock is proportional to the size of impulse responses to each of the shocks. Although inconsequential for the identified response of the target variable to the max-share shock, the convolution of shocks can substantially impact other quantities including the impulse responses of other variables and forecast error or dynamic variance decompositions. The rank one case illustrates a more general result that the principal eigenvector of Ξ will load on combinations of shocks with similarly shaped responses, emphasizing the role of the orthogonality condition in Theorem 1.

Two Shocks. Next, we consider the case with $N = 2$ and assume without loss of generality that $\Xi_{1,2} \geq 0$. Here, the solution to (9) implies:

$$\frac{\theta_1}{\theta_2} = \frac{\vartheta + \sqrt{\vartheta^2 + 4}}{2} \quad \text{where} \quad \vartheta \equiv \frac{\Xi_{1,1} - \Xi_{2,2}}{\Xi_{1,2}}. \tag{20}$$

The expression for θ_1/θ_2 is an increasing function of ϑ. The role of the orthogonality condition is captured by $\Xi_{1,2}$ in the denominator of ϑ. In particular, with $\Xi_{1,2} = 0$, the ratio θ_1/θ_2 tends to either zero or infinity. In contrast, $\Xi_{1,2} = \sqrt{\Xi_{1,1}\Xi_{2,2}}$ corresponds to the rank one case. The numerator of ϑ captures the relative size condition. Substantial weight will be placed on Shock 1 if $\Xi_{1,1} - \Xi_{2,2}$ is sufficiently large, i.e., the response to Shock 1 is sufficiently large relative to Shock 2 at the chosen horizons or frequencies.

[10] See Section 4.1.2 of Stock and Watson (2016) and references therein for further discussion of invertibility in VARs.

4 Implications and Interpretation

We now discuss the implications of the identification conditions in more detail. We show how they apply in the time and frequency domains, then discuss their interpretation more broadly.

4.1 Time Domain

Interpreting $\Xi^{time}(\mathcal{H})$. 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:

$$\Xi^{time}_{j,j'}(\mathcal{H}) = \psi_{\mathcal{H},j} \cdot \psi_{\mathcal{H},j'} = \begin{cases} \|\psi_{\mathcal{H},j}\|^2 = \sum_{h \in \mathcal{H}} \Psi^2_{h,ij} & \text{if } j = j' \\ \|\psi_{\mathcal{H},j}\| \, \|\psi_{\mathcal{H},j'}\| \cos \alpha_{jj'}(\mathcal{H}) & \text{if } j \neq j' \end{cases}, \tag{21}$$

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, 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 $\psi_{\mathcal{H},j'}$. $\Xi^{time}(\mathcal{H})$ thus 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'.

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

$$\psi_{\mathcal{H},1} \cdot \psi_{\mathcal{H},j} = 0 \text{ for all } j \neq 1 \tag{22}$$

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 violated in many situations. 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 almost trivially 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 (22) for any finite set of horizons, $\mathcal{H}$.

As we increase the number of shocks, N, condition (22) 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:

$$\|\psi_{\mathcal{H},1}\|^2 > \lambda_{\max}\left(\Xi^{time}_{2:N,2:N}(\mathcal{H})\right). \tag{23}$$

As discussed before, $\lambda_{\max}(\Xi^{time}_{2:N,2:N})$ depends on both magnitude and correlations of target variable responses to untargeted shocks $2,\dots,N$. 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 (21) makes clear that the measure of similarity is $\cos\alpha_{jj'}$.

An important property of Ξ^{time} reflected in (21) 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 $\mathcal{H} = \{0,\dots,H\}$, and $\|\psi_{\mathcal{H},j}\|$ depends on the entire response from impact through horizon H, and not only the relative size of individual responses at horizon H.[11] 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 $\psi_{\mathcal{H},j} \cdot \psi_{\mathcal{H},j'} = \|\psi_{\mathcal{H},j}\| \left\|\psi_{\mathcal{H},j'}\right\|$ or $\cos\alpha_{jj'} = 1$ for all (j,j') pairs. In other words, Ξ^{time} becomes rank one and we have solution (19), an extreme violation of the orthogonality condition. As a result, Dieppe et al. (2021) find mixed success in Monte Carlo simulations comparing non-accumulated max-share to the traditional max-share approach.

4.2 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:

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

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

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

$$\kappa_{i,j}(\omega) \equiv \sqrt{\left(\Gamma_{ij}^{\text{Re}}(\omega)\right)^2 + \left(\Gamma_{ij}^{\text{Im}}(\omega)\right)^2} \tag{24}$$

and *phase*:

$$\varphi_{i,j}(\omega) \equiv \tan^{-1}\left[-\frac{\Gamma_{ij}^{\text{Im}}(\omega)}{\Gamma_{ij}^{\text{Re}}(\omega)}\right] \tag{25}$$

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.[12]

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

$$\Xi_{j,j'}^{freq}(\Omega) = \begin{cases} \int_{\omega\in\Omega} \kappa_{i,j}^2(\omega)d\omega & \text{if } j = j' \\ \int_{\omega\in\Omega} \kappa_{i,j}(\omega)\kappa_{i,j'}(\omega)\cos\left(\varphi_{i,j}(\omega) - \varphi_{i,j'}(\omega)\right)d\omega & \text{if } j \neq j' \end{cases}, \tag{26}$$

which parallels the expression for $\Xi_{j,j'}^{time}$ in (21). For intuition, we first focus on the objects inside the integral, which are the elements of $\Xi_{j,j'}^{freq}(\Omega)$ when we focus on a singleton frequency band, $\Omega = \{\omega\}$. The squared gain, $\kappa_{i,j}^2$ 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:

$$\int_{\omega\in\Omega} \kappa_{1,1}(\omega)\kappa_{1,j'}(\omega)\cos\left(\varphi_{1,1}(\omega) - \varphi_{1,j'}(\omega)\right)d\omega = 0 \ \text{ for all } \ j \neq 1, \tag{27}$$

and the relative size condition can be written:

$$\int_{\omega\in\Omega} \kappa_{1,1}^2(\omega)d\omega > \lambda_{\max}\left(\Xi^{freq}(\Omega)_{2:N,2:N}\right). \tag{28}$$

These expressions, together with (24)–(26), make a straightforward interpretation in the frequency domain difficult. 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 in Angeletos et al. (2020). More generally, (26) requires integrating nonlinear functions of the time-

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

domain representation of the impulse responses over the frequency band, Ω. Consequently, it is arguably a 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. This discrepancy helps explain the contrasting results that Angeletos et al. (2020) find when comparing max-share in the time and frequency domains.

4.3 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. We specifically emphasize the amount of prior knowledge the econometrician requires to use max-share as a plausible identification approach.

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 then doing a Cholesky decomposition of Σ so that C is lower triangular. It also corresponds to the rank one case described in Section 3.5. The internal instrument literature places zero restrictions on the $(1, j)$ elements on the top row of C for $j > 1$. In the max-share setting, this is 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 0 on impact. Otherwise, the weights on the true shocks are given by (19).

When we instead target the FEV for horizon $H > 0$, the restrictions are no longer restricted to the response on impact. Instead, they now apply to the entire response over horizons $0, \ldots, H$. The zero restrictions are then replaced by the assumption that the remaining impulse responses lie on the null space of $\psi_{0:H,1}$. This is arguably more difficult to satisfy in many models. Specifically, 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 unless the untargeted responses happen to be orthogonal to the targeted one.

Sign Restrictions. While max-share does not directly impose sign restrictions, it can still implicitly impose some 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 restriction where the assumptions are on the response of only the target shock (Uhlig, 2005; Arias et al., 2018), the max-share identification conditions apply to both the targeted and untargeted shocks jointly. In other words, one needs to be prepared 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 Illustrative Examples

As an illustration, we consider the general form for a log-linearized dynamic stochastic general equilibrium (DSGE) model:

$$x_t = Ax_{t-1} + B\varepsilon_t \tag{29}$$

$$y_t = Cx_{t-1} + D\varepsilon_t, \tag{30}$$

where $\varepsilon_t \sim \mathcal{N}(0, I)$. We can then write:

$$\Psi_{i,j,h} = \begin{cases} D_{i,j} & h = 0 \\ C_{i,1:N_x} A^{h-1} B_{1:N_x,j} & h \geq 1 \end{cases}, \tag{31}$$

where N_x is the dimensionality of the state vector x_t. 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.[13]

Differences in the dynamic responses to different shocks beyond horizon 0 thus depend on differences in exposures, $B_{1:N_x,j}$, of the states x_t to the shocks and dynamics of those states, as captured by A. If the states have similar autoregressive properties or exposures to the structural shocks, then a given variable will have a similar dynamic response to the different shocks. Consequently, max-share will identify a convolution of these shocks.

For intuition, consider the case with $N_x = 1$ so that x_t and A are scalars. This is the case in a standard real business cycle model (King, Plosser, and Rebelo, 1988) with iid shocks, where

[13] See Fernández-Villaverde, Rubio-Ramírez, Sargent, and Watson (2007) for conditions under which a VAR can reproduce the model impulse responses.

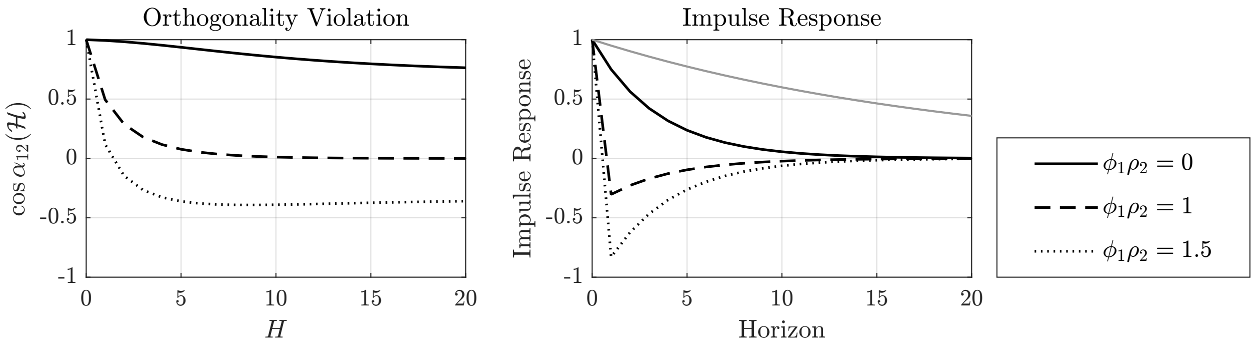

Figure 2: Orthogonality violations in the time domain for ARMA(1,1) impulse responses. **Left:** Orthogonality violations, $\cos\alpha_{12}(\mathcal{H})$; **Right:** Impulse responses. Black lines correspond to Shock 1; gray line corresponds to Shock 2.

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 A. Suppose again without loss of generality that the target shock is ordered first. Then, for a given horizon or frequency band and $D_{i,1}$, we require a specific $D_{i,j}/B_{i,j}$ for all $j \neq 1$ in order to satisfy orthogonality.

5.1 ARMA(1,1) Impulse Responses

We now study the orthogonality condition and deviations from it in (29)-(30) with $A = \text{diag}\{\rho_j\}$, $B = \text{diag}\{\rho_j - \phi_j\}$ with $|\rho_j| < 1$ for $j \in \{1, 2\}$, and $C = D = (1, 1)'$. The impulse responses would then resemble 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:

$$\psi_{j,h} = \begin{cases} 1 & h = 0 \\ (\rho_j - \phi_j)\rho_j^{h-1} & h > 0 \end{cases}. \tag{32}$$

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

$$\psi_{1,\mathcal{H}} \cdot \psi_{2,\mathcal{H}} = 1 + (\rho_1 - \rho_1)(\rho_2 - \phi_2)\frac{1 - (\rho_1\rho_2)^H}{1 - \rho_1\rho_2} \tag{33}$$

and orthogonality is achieved when:

$$H = \log\left(1 + \frac{1 - \rho_1\rho_2}{(\rho_1 - \phi_1)(\rho_2 - \phi_2)}\right) / \log(\rho_1\rho_2). \tag{34}$$

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 need 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

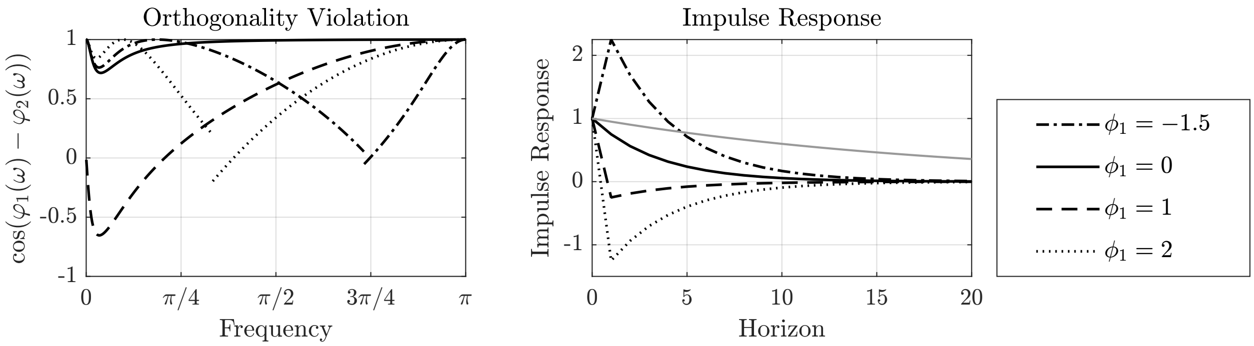

Figure 3: Orthogonality violations in the frequency domain for ARMA(1,1) impulse responses. **Left:** Orthogonality violations, $\cos(\varphi_1(\omega) - \varphi_2(\omega))$; **Right:** Impulse responses. Black lines correspond to Shock 1; gray line corresponds to Shock 2.

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 $H \to \infty$. However, this is no longer the case once we increase ϕ_1 further such that $\phi_1 \rho_2 = 1.5$. In this case, we again have relatively strong violations of orthogonality except around the solution to (34) of 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, one can construct examples in which max-share approximately or even exactly identifies the true shock of interest, but 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 (32). To simplify our analysis, we focus on the single frequency case with $\Omega = \{\omega\}$.

Defining:

$$\varphi_j^{\mathsf{x}}(\omega) = -\tan^{-1}\left[\frac{\mathsf{x}_j \sin \omega}{1 - \mathsf{x}_j \cos \omega}\right] \quad \text{for } \mathsf{x} \in \{\rho, \phi\}, \tag{35}$$

we can write the phase difference as:

$$\varphi_1(\omega) - \varphi_2(\omega) = (\varphi_1^{\rho}(\omega) - \varphi_1^{\phi}(\omega)) - (\varphi_2^{\rho}(\omega) - \varphi_2^{\phi}(\omega)). \tag{36}$$

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 (36) 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 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 the left panel 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 as to 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 or even on the specific frequency to target.

5.2 Supply and Demand

Another instructive case takes A and B to be diagonal and $D = CB$, so that x_t consists of independent AR(1) processes and the model can be written as:

$$y_t = Cx_t = CAC^{-1}y_{t-1} + D\varepsilon_t. \tag{37}$$

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

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

$$q_t^d = -\gamma^d p_t + \eta_t^d, \tag{39}$$

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:

$$\eta_t^{\times} = \rho^{\times}\eta_{t-1}^{\times} + \sigma^{\times}\varepsilon_t^{\times}, \tag{40}$$

where $\varepsilon_t^{\times} \overset{iid}{\sim} \mathcal{N}(0,1)$ with $\times \in \{s, d\}$. 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 (37). 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 (38)–(39) is static.[14] Details are in the Online Appendix.

Time Domain: Technology Shock. As an illustration in the time domain, we identify the supply shock using max-share identification on output 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

[14]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.

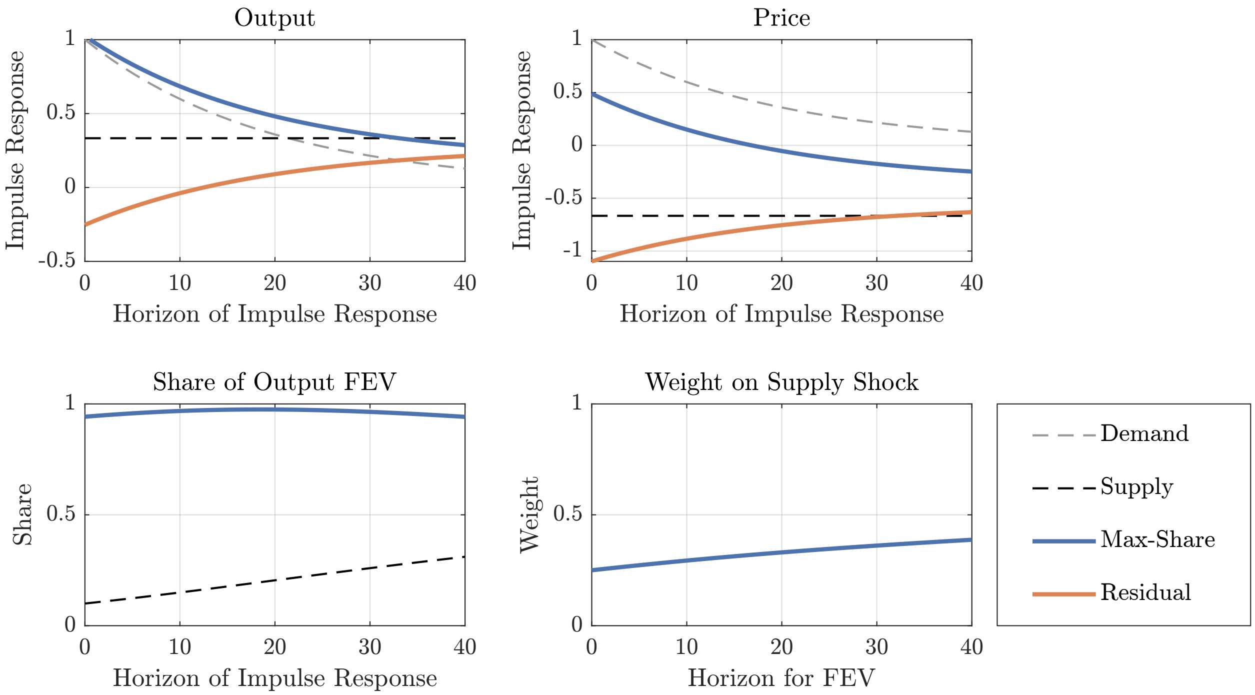

Figure 4: Permanent technology shock in supply and demand example identified via max-share in the time domain. **Top panels:** True and identified impulse responses; **Bottom left:** Contribution to output FEV; **Bottom right:** Weight that max-share identified technology shock places on true supply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

identification (targeting labor productivity instead of output). We consider the parameter values:

$$\gamma^s = 1.00, \qquad \rho^s = 1.00, \qquad \sigma^s = 1.00$$
$$\gamma^d = 0.50, \qquad \rho^d = 0.95, \qquad \sigma^d = 1.50$$

The supply shock is a random walk and thus the only shock with a permanent effect on q_t. The demand shock persistence $\rho^d = 0.95$ is less persistent than the choice of 0.98 used by Francis et al. (2014). Following Francis et al. (2014), we set $\mathcal{H} = \{0, \ldots, 40\}$, targeting horizon 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, the lower left panel of 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, the top two panels show that the max-share shock produces a positive response in both output and price, resembling a demand shock. Next, the bottom left panel 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 shock only places a weight of 0.39 on the supply shock, as seen in the bottom right panel. 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 thus imposes restrictions on the untargeted responses that can potentially imply counterintuitive results.

Frequency Domain: Business Cycle Shock. In the frequency domain, we follow Angeletos et al. (2020) and target the response of output at frequency band $\Omega = \left[\frac{2\pi}{32}, \frac{2\pi}{6}\right]$. They label this as the "main business cycle" shock, finding the striking result that the identified shock produces a large responses in real variables but a small response in inflation.

First, we consider the case with symmetric processes for the supply and demand shocks:

$$\gamma^s = \gamma^d = 1.00, \qquad \rho^s = \rho^d = 0.95, \qquad \sigma^s = \sigma^d = 1.00.$$

The top panel of Figure 5 shows that the max-share shock qualitatively resembles the main business cycle shock in Angeletos et al. (2020), producing a positive response in output, q_t, but no response in price, p_t. The residual shock displays the opposite behavior, producing a zero response in output but a substantial response in price. In this model, the lack of response in output to the residual shock occurs as long as the two true underlying shocks have the same persistence, $\rho^s = \rho^d$. Since this implies that the response of output to both the max-share and residual shocks must have AR(1) dynamics with persistence $\rho^s = \rho^d$, the residual shock is forced to produce a zero response. The responses of price then depend on the elasticities, $\{\gamma^s, \gamma^d\}$, and standard deviations, $\{\sigma^s, \sigma^d\}$. In this context, valid max-share identification rests on the argument that the one other shock in the economy affects prices but not quantities.

We deviate from this knife-edge case by setting $\rho^s \neq \rho^d$. The lower panel of Figure 5 shows results when we choose $\rho^s = 0.85$ but keep all other parameters unchanged. The max-share shock now produces a positive but relatively small response in price. However, the residual shock also generates impulse responses in output and price that have the same sign. In other words, with this parameterization, max-share identification implies that both the supposed main business cycle shock of Angeletos et al. (2020) and the residual shock have the features of demand shocks, essentially ruling out the presence of supply shocks.

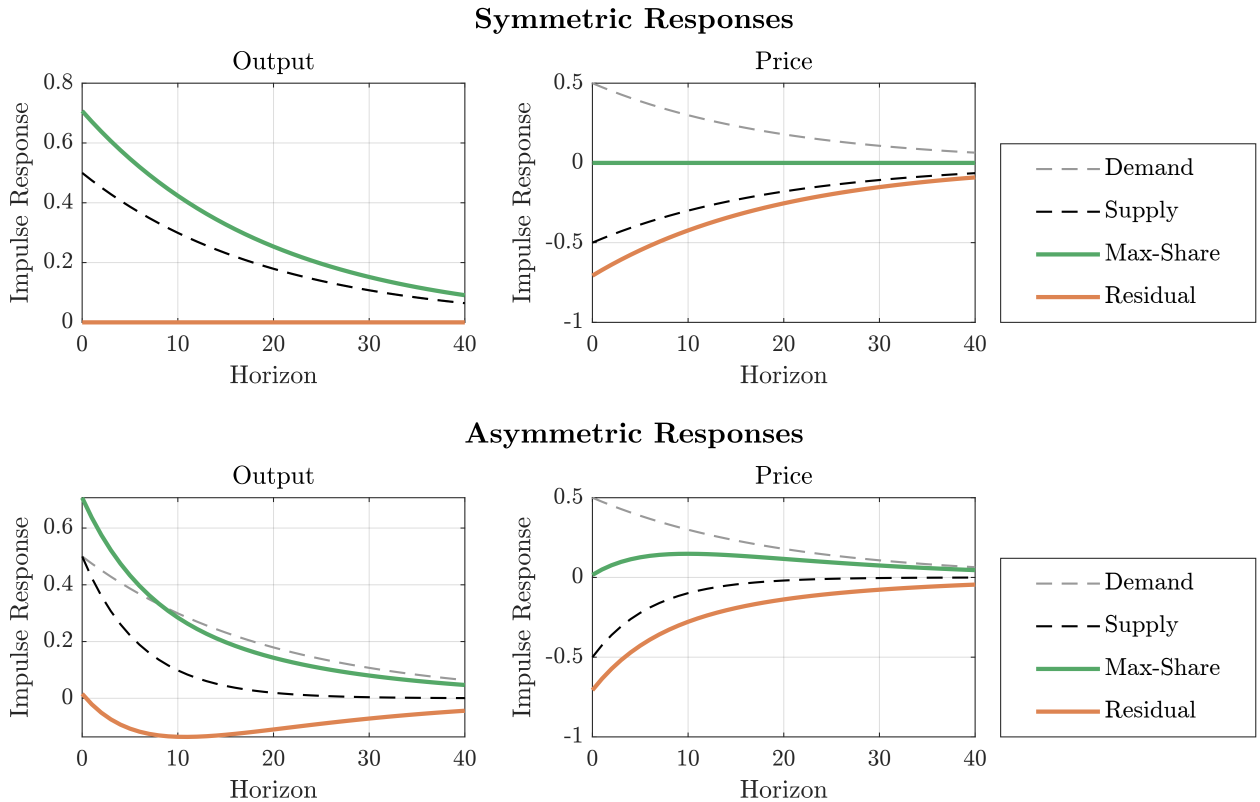

Figure 5: Main business cycle shock in supply and demand example identified via max-share in the frequency domain. **Top panels:** Symmetric responses with $\rho^s = \rho^d = 0.95$; **Bottom panels:** Asymmetric responses with $\rho^s = 0.85$ and $\rho^d = 0.95$. Dashed lines indicate true responses; green and orange solid lines correspond to identified max-share and residual shocks, respectively.

This message echoes our findings from the time domain—in using max-share identification, researchers need to be careful of implications for not only the targeted shock, but also the untargeted ones. In both cases discussed here, the orthogonality conditions lead to untargeted shocks with responses that seem unlikely from economic theory.

6 Empirical Application

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 also 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.

6.1 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.[15]

Building on a large literature (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 only 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 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 $\mathcal{H} = \{0, \dots, 40\}$. That is, we find the shock that contributes the most to the FEV of TFP over horizon $H = 40$ quarters. 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 it identifies the surprise shock as being the one that explains all residual variation in TFP after controlling for lagged variables.

Identification of the business cycle shock is based on seminal work by Angeletos et al. (2020). We closely follow their approach and use max-share in the frequency domain, whereby we target GDP at frequency band $\Omega = \left[\frac{2\pi}{32}, \frac{2\pi}{6}\right]$.[16] Angeletos et al. (2020) argue that this "main business cycle shock" resembles a demand shock that has minimal effect on inflation. They rule out the role of supply shocks since TFP has a relatively small response to their business cycle shock. 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. The TFP news and surprise

[15] 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.

[16] 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.

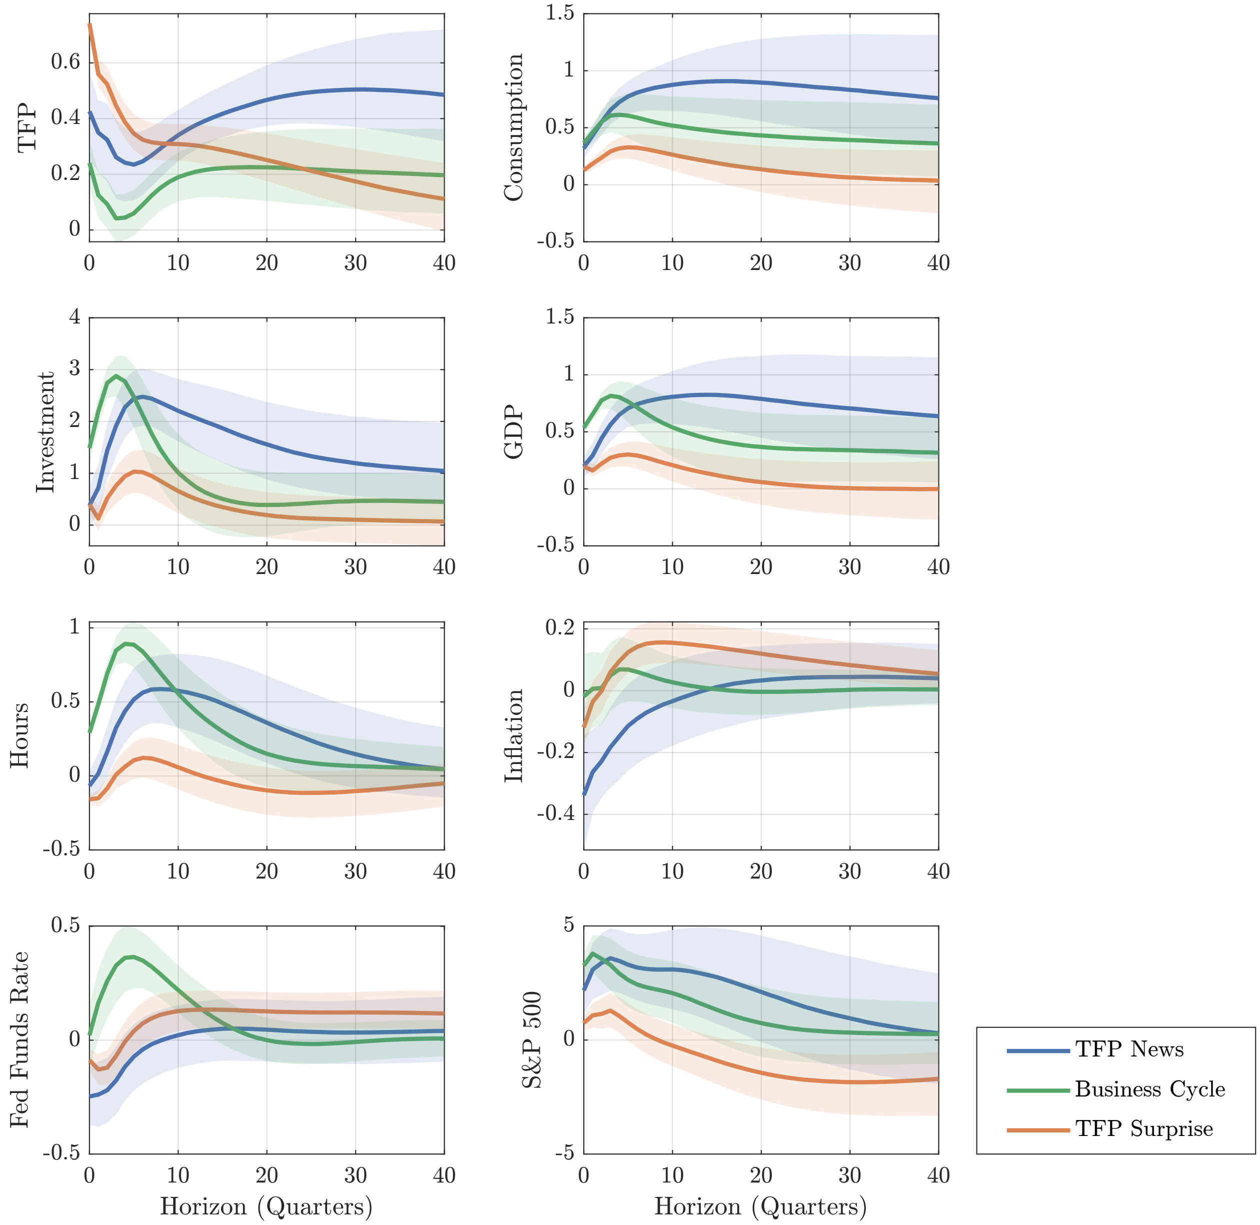

Figure 6: Posterior estimates for impulse responses to identified shocks. **Solid lines:** Median response; **Shaded regions:** 68% error bands.

shocks are 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.

6.2 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 shocks are likely contaminated.

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 a hockey stick shape, with a substantial response at short horizons.[17]

While the analysis for the business cycle shock is less accessible because it is identified in the frequency domain, we observe that it produces highly persistent response in GDP that is mirrored by the TFP news shock. While the shape of the responses over the short 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.

The contributions of the shocks to the FEV of each variable, shown in Figure 7, provide further suggestive evidence that the shocks are not all be cleanly identified. The contributions of the TFP news and surprise shocks to TFP sum to more than one around horizon 20 and the contributions of the business cycle and TFP news shocks to GDP sum to more than one around horizon 12 at their medians.

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

$$\frac{\hat{\beta}}{\sqrt{1 - \hat{\beta}^2}} \tag{41}$$

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 an 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 faction of the weight max-share places on the (combination of) shock(s) generating $\hat{\psi}$, we have that (41) will be $\frac{\zeta}{1-\zeta}$. In particular, if max-share does isolate the true structural shock of interest ($\zeta = 0$), then (41) will be zero. If half the max-share weight is on the (combination of) shock(s) generating $\hat{\psi}$ ($\zeta = 0.5$), then (41) will be one.

[17] 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).

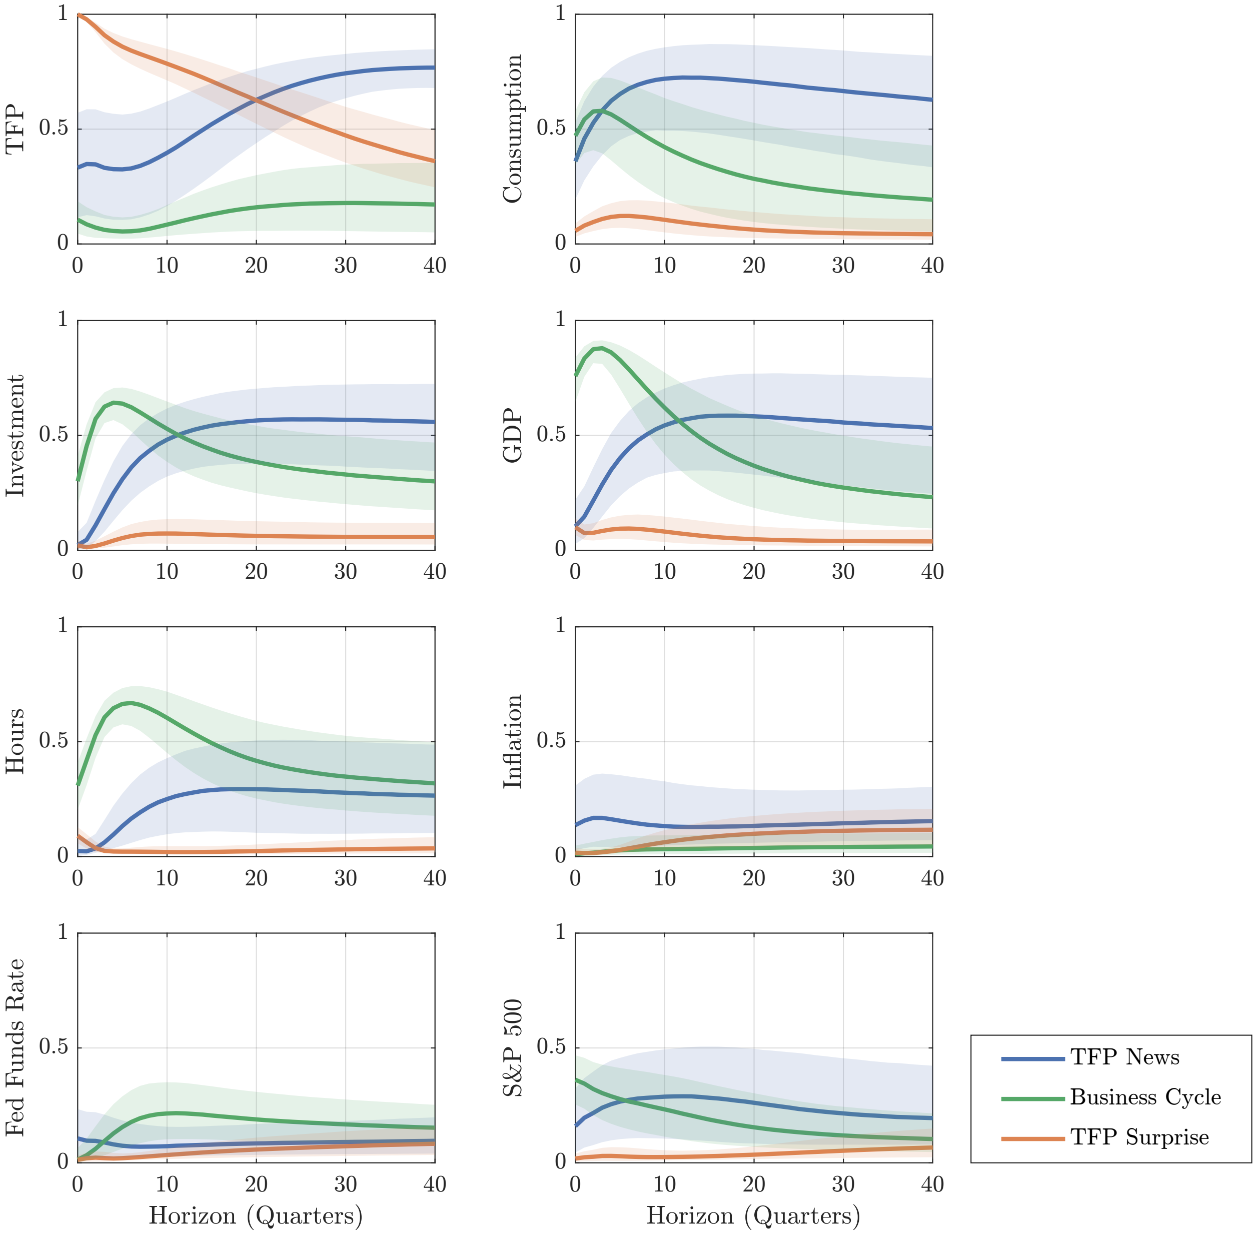

Figure 7: Posterior estimates for FEV contributions of identified shocks. **Solid lines:** Median response; **Shaded regions:** 68% error bands.

Figure 8 plots the posterior for (41) and shows that 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 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 has centered only on the responses of the target variables, TFP and GDP. This is motivated by our theoretical results and the primitive problems, (4) and (6), none of which

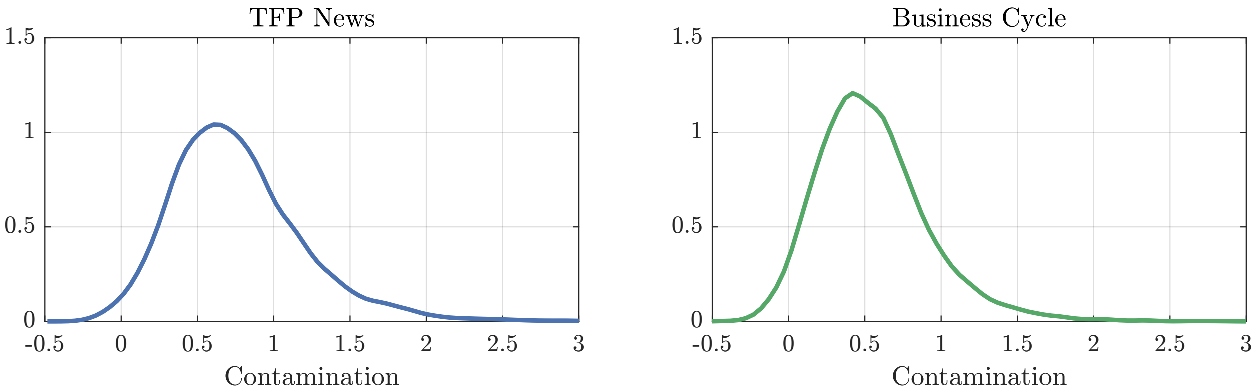

Figure 8: Posterior distribution for contamination of TFP news (left) and business cycle (right) max-share shocks by identified TFP surprise and TFP news shocks, respectively.

reference the untargeted variables. In Figure 6, there are some 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 by 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. This contradicts the claim by Angeletos et al. (2020) that small contributions to TFP and inflation by the main business cycle shock implies that we should rule out a role for shocks that affect TFP or inflation in driving the business cycle. More generally, the responses of untargeted variables alone are insufficient to argue for the validity of max-share, as they do not directly correspond to the conditions in Theorem 1.

7 Conclusion

In many applications max-share is used as a seemingly straightforward way to identify structural shocks in VARs that avoids problems associated with other identification approaches, especially when considering non-standard shocks or at longer horizons. We show in this paper that, in fact, max-share identification is quite problematic and that stringent conditions are required to accurately identify the shock of interest. First, the impulse response of the target variable to the shock of interest must be orthogonal to the corresponding responses to all other shocks. We show in a wide range of examples not only that this is hard to satisfy, but that violations can lead to max-share severely misidentifying the shock of interest. Second, the target response must be large relative to combinations of responses to the untargeted shocks in the VAR, as summarized by the corresponding eigenvalues. This contrasts with a claim that a shock can be identified through max-share if it produces a larger response in the target variable than each of the other shocks individually. Given these concerns, we develop a diagnostic for the degree of contamination in the max-share shock. It relies on identifying an individual or combination of untargeted shocks which is straightforward to implement in a VAR.

Although the results in this paper raise serious concerns with the use of max-share identification to isolate a shock of interest, it does not preclude the methodology's usefulness. It can still be useful for obtaining bounds on impulse responses, as was the original goal of Faust (1998). In addition, the max-share impulse responses can be used for indirect inference, e.g., as a target for a simulated method of moments estimation in a structural economic model.

References

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.

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.

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.

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.

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.

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.

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.