# 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

September 19, 2024

**Abstract**

Max-share identification relies on a decomposition of the forecast error variance (FEV) over a target horizon. Consequently, it often conflates multiple shocks because the contribution to the FEV depends on the impulse responses at untargeted horizons and the shapes of the responses to untargeted shocks. We alleviate the issues using a so-called "single horizon" alternative that focuses narrowly on the actual target horizon. We characterize the identified shock in terms of true structural shocks in the single horizon problem and show that this typically bounds results in the literature's usual implementation. Using a numerical demand and supply example and an empirical news shock application, we show that the traditional max-share approach inadvertently places weight on untargeted transitory shocks, a problem that the single horizon approach avoids.

*We are grateful to participants at 2024 Annual Conference of the Society for Computational Economics 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

Structural vector autoregressions (SVARs) are the macroeconomist's dominant tool for the empirical study of economic policies and the sources of business cycles. Key to this framework is the ability to identify structural shocks, that is, the mutually independent underlying drivers of economic activity. A variety of alternative approaches have proliferated since early approaches to structural identification, such as recursive orderings of shocks (Sims, 1980). Recently, the so-called max-share approach has become increasingly popular. 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 or frequency.

A prominent example is the identification of a news shock affecting total factor productivity (TFP). A TFP news shock is the expectation of a future change in TFP even though it has not yet been realized. Even before the actual change in TFP, the news shock has an effect on economic activity as households and firms respond to their expectations. Consequently, the literature has sought to identify structural TFP news shocks by maximizing their contribution to the FEV of TFP over some medium or long horizon (Beaudry and Portier, 2006; Barsky and Sims, 2011).

We show in this paper that structural identification via max-share, as it has been employed in the applied literature, is problematic since it conflates the effects of a multitude of shocks. As a result, the literature misstates both the impulse responses and the importance of various shocks. This problem is pervasive, which we demonstrate using simple analytical and numerical examples as well as an application to the news shock literature. Despite the widespread use of max-share, we argue that the literature lacks a general analysis of its performance and validity. This paper fills that gap by discussing what determines the weights that max-share identification places on the true structural shocks of a general VAR setup and the tension between these forces and the usual justification for max-share. Finally, we propose a closely related alternative that overcomes key issues we highlight and makes the underlying identification assumptions more transparent.

**Max-Share and Its Problems.** To understand the identification problems that can arise with the max-share approach, consider the problem for some horizon $H$ and target variable $i$, as originally proposed by Faust (1998) and followed by much of the subsequent literature. Max-share identifies a shock by finding the rotation that maximizes the identified shock's

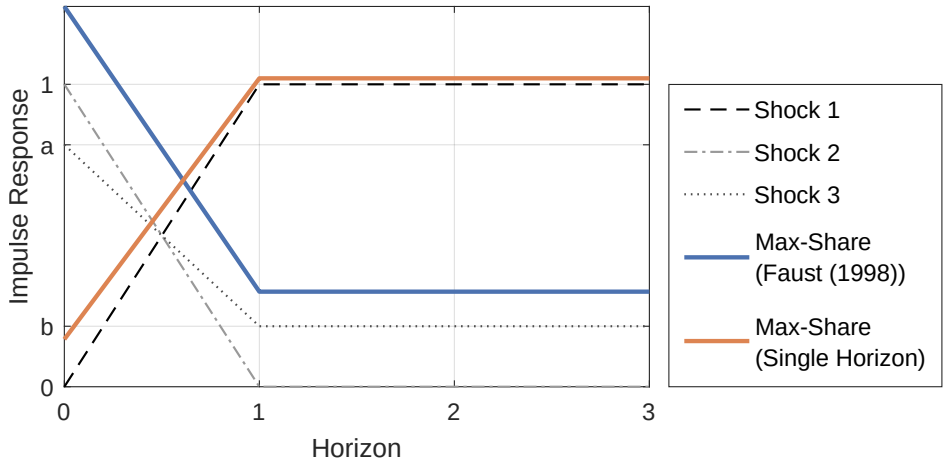


Figure 1: Impulse responses to shocks in stylized example. Discussion primarily considers general choice of $a$ and $b$; figure plots case with $a = 0.8$ and $b = 0.2$.

contribution to the FEV of the target variable $i$, which is the $i$th diagonal element of:

$$\sum_{h=0}^{H} \Psi_h \Psi_h', \tag{1}$$

where $\Psi_h$ is a matrix whose $(i, j)$ entry is the response of variable $i$ to shock $j$ at horizon $h$. One can think of this as choosing a set of weights to attach to the shocks, with the true shock of interest being correctly identified if all the weight is placed on it.

By focusing on the FEV, max-share conflates shocks in a way that contradicts its underlying assumptions. A common motivation for max-share is that the shock of interest dominates at the target horizon $H$ in the sense that its impulse response is substantially larger than those corresponding to the other shocks at the target horizon (Francis, Owyang, Roush, and DiCecio, 2014; Barsky and Sims, 2011). For instance, an identified technology shock preferably has the largest impact on a technology variable, such as TFP. However, the summation expression (1) makes clear that the FEV depends not only on horizon $H$ but also on all prior horizons $h < H$. Moreover, the matrix product in the summation causes the weights to depend on relationships between impulse responses to different shocks that the econometrician is unlikely to have prior knowledge about. These two pitfalls potentially lead to highly misleading identification.

As a simple example, consider a structural VAR in which the impulse responses of the target variable to the three true shocks are shown in Figure 1. Suppose we seek to identify shock 1 and consider a target horizon $H = 1$, with the intuition that it should have a relatively large response after impact. Max-share identifies a shock that is a weighted average of the

three true shocks. Denoting the weight on shock $j$ by $\theta_j$, we can show:

$$\frac{\theta_1}{\theta_2} = \frac{b}{a} \quad \text{and} \quad \frac{\theta_1}{\theta_3} = \frac{b}{a^2 + b^2}. \tag{2}$$

These relative weights highlight the two issues with max-share identification described above.

First, the weights depend not only on the impulse responses at the target horizon $h = 1$, but also at horizon $h = 0$. The weight on shock 1 is decreasing in the impact response, $a$, to shock 3, even though it is irrelevant to the size of the impulse response at the target horizon. These dependencies arise even though in many cases the econometrician may not intend to impose any beliefs about the shorter horizons $h < H$.

Second, the weight on the shocks depends on the relative shapes of the impulse responses to different shocks (e.g., how much they decay with horizon or at what horizon they peak). When $|b| < |a|$, reducing the magnitude of $b$ increases the weight on shock 3 relative to shock 1 despite reducing the size of the impulse response to shock 3. In the limit, when $b = 0$, zero weight is placed on the target shock 1 even though it is the only shock with a non-zero response at the target horizon. This arises because $b = 0$ implies that the responses to shocks 2 and 3 are identical up to scale, increasing their joint contribution to the FEV. In other words, to justify the use of max-share, the econometrician needs to impose beliefs not only on the relative magnitude of the response to the shock of interest, but also the shapes of the impulse responses for all the other shocks.

**A Simple Alternative.** As a safeguard against this misidentification issue, we propose that researchers focus on the impulse response at a specific horizon rather than the FEV, i.e., concentrate on a single term in the summation (1). In doing so, we remove dependence on all horizons $h \neq H$ and make the exact shapes of the impulse responses irrelevant. We refer to this approach as the *single horizon* problem.

The solution to the single horizon problem yields relative weights on the true shocks that are exactly the relative size of their impulse responses at the target horizon. As with any identification scheme, the single horizon problem is not always appropriate, but our results clarify the necessary conditions for it to work well. Moreover, we show that the weights for the Faust (1998) max-share approach are typically bounded by the weights obtained from the single horizon problems for horizons 0 to $H$. Therefore, the single horizon problem for an appropriately chosen $h$ performs as well as or better than the FEV max-share identification.

In the stylized example above, the single horizon problem yields:

$$\theta_1 = \frac{1}{\sqrt{1 + b^2}}, \quad \theta_2 = 0, \quad \theta_3 = \frac{b}{\sqrt{1 + b^2}}. \tag{3}$$

The weight on shock 1 increases as $|b|$ decreases, with $\theta_1 = 1$ when $b = 0$. This intuitive relationship is in stark contrast to the Faust (1998) approach above, where decreasing $b$ can in fact decrease the weight, $\theta_1$, on shock 1 even though it increases the relative size of the impulse response to shock 1. In addition, no weight is placed on shock 2 since its corresponding response at horizon 1 is zero, as is intended. While the weight on shock 1 will be small if $b$ is large, it is likely easier to impose assumptions on $b$ than on the joint distribution of $a$ and $b$. Indeed, Figure 1 shows that with $a = 0.8$ and $b = 0.2$, the Faust (1998) approach yields an identified impulse response that is completely different than the true response, whereas the single horizon impulse response closely resembles the true one.

**Related Literature and Max-Share Applications.** The idea of structural identification via the FEV goes back to 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). In a well-known contribution Francis et al. (2014) apply this approach to the issue of long-run identification, that is, the identification of shocks that have permanent effects. The original approach of Blanchard and Quah (1989) and popularized by Gali (1999) is based on the idea of putting identifying restrictions directly on the long-run variance, which is tenuous in small samples. In contrast, Francis et al. (2014) suggest a max-share approach, identifying a permanent supply shock as the one that maximizes its share in the FEV of labor productivity at a long horizon. We consider a simple 2-variable demand and supply system in Section 3 to illustrate the potential pitfalls of max-share identification in such settings.

Kurmann and Sims (2021) recognized the usefulness of max-share for the identification of news shocks. They implement a strategy whereby a TFP news shock is identified as the shock that maximizes the FEV of observed TFP over a 10-year horizon. This time span is chosen based on a priori reasoning about the time taken for innovation to diffuse and for inventions to become economically viable. Their approach has become the standard empirical framework for identifying and estimating the effects of news shocks. Related work by Barsky and Sims (2011) and Ben Zeev and Khan (2015) maximize the contribution to the sum of FEVs over various horizons subject to additional constraints. The literature has produced a wide range of estimates for the contribution of these identified shocks to business cycle frequency output fluctuations, as summarized in Ramey (2016). We revisit this prominent empirical application and contrast results from the standard max-share approach in the literature with our single horizon problem.

In addition to the news shock literature, the max-share approach has been used to identify uncertainty shocks (Caldara, Fuentes-Albero, Gilchrist, and Zakrajšek, 2016), credit

shocks (Mumtaz, Pinter, and Theodoridis, 2018), business cycle shocks (Angeletos, Collard, and Dellas, 2020), sentiment shocks (Levchenko and Pandalai-Nayar, 2020), risk premium shocks (Basu, Candian, Chahrour, and Valchev, 2023), and exchange rate shocks (Chahrour, Cormun, De Leo, Guerrón-Quintana, and Valchev, 2024). For many of these shocks, it is a challenge to put sufficient zero and sign restrictions (Sims, 1980; Uhlig, 2005; Arias, Rubio-Ramírez, and Waggoner, 2018) or to find suitable instruments (Mertens and Ravn, 2013; Stock and Watson, 2018) for identification. An attraction of max-share identification is that it only requires the seemingly simple assumption that the target shock is important for a particular variable at some horizon or frequency. Our results reveal that this assumption 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. 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 by Kurmann and Sims (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 separately identify news and uncertainty shocks, max-share yields shocks that are highly correlated instead of being independent, an issue Carriero and Volpicella (2024) resolves by jointly identifying multiple shocks by max-share. We build on these papers by presenting more general results, which we then connect back to specific applications through numerical or empirical examples.

**Outline.** Section 2 describes the setup and general results. Section 3 illustrates the results numerically through a bivariate supply and demand example. Section 4 presents an empirical news shock application. Finally, Section 5 concludes.

# 2 Properties of Max-Share

## 2.1 Setup

Consider a general structural VAR:

$$Y_t = \sum_{\ell=1}^{L} B_\ell Y_{t-\ell} + C\varepsilon_t, \tag{4}$$

where $Y_t$ is an $N \times 1$ vector and $\varepsilon_t$ is i.i.d. over time, with $E[\varepsilon_t] = 0$ and $E[\varepsilon_t \varepsilon_t'] = I$. We can write the moving average representation:

$$Y_t = \sum_{h=0}^{\infty} \Psi_h \varepsilon_{t-h} \tag{5}$$

where the $N \times N$ matrix $\Psi_h$ summarizes the impulse responses at horizon $h$.[1] Each column of $\Psi_h$ 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 $\Psi_h$.

To identify $\Psi_h$, we consider the problem:

$$\arg\max_{\theta} \frac{\delta_i' \left[\sum_{h \in \mathcal{H}} \widetilde{\Psi}_h \theta \theta' \widetilde{\Psi}_h'\right] \delta_i}{\delta_i' \left[\sum_{h \in \mathcal{H}} \widetilde{\Psi}_h \widetilde{\Psi}_h'\right] \delta_i} \quad \text{subject to} \quad \theta'\theta = 1, \tag{6}$$

where $\theta$ is the rotation or vector of weights that we are solving for, $\delta_i$ is a vector with 1 in the $i$th entry and 0 everywhere else, and $\mathcal{H}$ is set of horizons chosen by the econometrician based on economic theory. Finally, $\widetilde{\Psi}_h$ is an arbitrary rotation of the structural shocks satisfying $\widetilde{\Psi}_h \widetilde{\Psi}_h' = \Psi_h \Psi_h'$. For exposition, we will take $\widetilde{\Psi}_h = \Psi_h$, so that we can interpret the $j$th entry of the solution to (6) as the weight that the max-share shock places on the $j$th true shock.[2]

We divide the possible choices of $\mathcal{H}$ into two cases:

- **Single horizon:** $\mathcal{H} = \{h\}$ is a singleton.
- **Multiple horizon:** $\mathcal{H} = \{h_0, \ldots, h_H\}$ consists of more than one horizon.

The typical use of (6) in max-share identification, following Faust (1998), corresponds to the multiple horizon case with $\mathcal{H} = \{0, 1, \ldots, H\}$. This maximizes the contribution of the identified shock to the FEV of variable $i$. It 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 internal instrument identification (Noh, 2018; Plagborg-Møller and Wolf, 2021).

[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 (4) and $(\mathbb{B}^h)_{1:N,1:N}$ denotes the upper-left $N \times N$ submatrix of $\mathbb{B}^h$.

[2] This choice of $\widetilde{\Psi}_h$ affects the solution for $\theta$ in (6), but does not change the implied impulse responses to the max-share shock. In practice, since the true responses are unknown, a common choice is to take $\widetilde{\Psi}_h$ to be the lower triangular matrix from the Cholesky decomposition of $\Sigma$.

We will show that the single horizon approach provides improved performance and transparency over the multiple horizon problem in identifying the true shock. Nevertheless, equation (6) also makes clear that the two approaches are equally straightforward to implement.

## 2.2 Theoretical Results and Simulation Evidence

We now establish several general results on the solution to (6). The closer the weight on the true shock is to one, the better max-share identification approximates the true shock. How this translates to individual impulse responses resembling the true responses depends both on the weights and on those responses. However, since we will show that (6) only depends on the impulse responses to the target variable, we only present results on the weights and not the impulse responses. Sections 3 and 4 will discuss how the impulse responses could be affected in practice. All proofs are provided in the Appendix.

**Equivalent Eigenproblem.** It is useful to first recast the problem as an eigenproblem, as first discussed in Faust (1998).

**Lemma 1.** *Solving (6) is equivalent to solving the eigenproblem for*

$$\sum_{h\in\mathcal{H}} \widetilde{\Psi}_h' \delta_i \delta_i' \widetilde{\Psi}_h \tag{7}$$

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

Taking $\widetilde{\Psi}_h = \Psi_h$, Lemma 1 provides intuition for the weights that max-share identification places on the true shocks. The expression (7) is an $N \times N$ matrix whose $(j, j')$ entry is the dot product

$$\psi_{\mathcal{H},j} \cdot \psi_{\mathcal{H},j'}, \text{ where } \psi_{\mathcal{H},j} \equiv \begin{bmatrix} \Psi_{h_0,ij} & \cdots & \Psi_{h_H,ij} \end{bmatrix}' \tag{8}$$

and the dependence on $i$ is suppressed to economize on notation. In particular, we have:

$$\psi_{\mathcal{H},j} \cdot \psi_{\mathcal{H},j'} = \begin{cases} \|\psi_{\mathcal{H},j}\|^2 = \sum_{h\in\mathcal{H}} \Psi_{h,ij}^2 & \text{if } j = j' \\ \|\psi_{\mathcal{H},j}\| \, \|\psi_{\mathcal{H},j'}\| \cos\alpha_{jj'} & \text{if } j \neq j' \end{cases}, \tag{9}$$

where we omit the dependence of $\alpha_{jj'}$ on $\mathcal{H}$ and $i$ to reduce notation. The $j$th diagonal element of (7) captures the size of the $j$th true impulse response, which is also the contribution of shock $j$ to the FEV of variable $i$. The $(j, j')$ off-diagonal element captures the similarity of the $j$ and $j'$ impulse responses, scaled by their sizes. The angle $\alpha_{jj'}$ captures how close the vectors are. For instance, $\cos\alpha_{jj'} = 1$ implies that the vectors are parallel (i.e., the impulse

response of variable $i$ to the $j$th and $j'$th are identical up to scale) and $\cos\alpha_{jj'} = 0$ implies that the vectors are orthogonal. In the example in the Introduction, $b = 0$ corresponds to $\cos\alpha_{13} = 0$ since the impulse response to shock 3 is zero whenever the response to shock 1 is non-zero and $\cos\alpha_{23} = 1$ since the impulse responses to shock 3 is identical to that to shock 2, scaled by $a$. Intuitively, one can think of the max-share problem (6) as being analogous to solving for the principal component of a set of variables (without normalizing them to have unit variance), with the impulse response for variable $i$ to each shock here corresponding to a variable in the principal component analysis context. The comparison suggests two forces at play in the max-share problem, both of which can contribute to misidentification.

First, a shock will receive greater weight if its corresponding impulse response is larger. Importantly, this magnitude depends equally on the size of the impulse response at all horizons in $\mathcal{H}$. Thus, in the typical multiple horizon problem with $\mathcal{H} = \{0, \ldots, H\}$, we are implicitly assuming a large response of variable $i$ to the true shock of interest not only at horizon $H$, but also at horizons prior to that. However, we often have little prior knowledge of the relative size of shocks at short horizons even if we believe that one shock is likely to dominate at horizon $H$. The multiple horizon problem then implicitly imposes stronger assumptions than one might intend. In the example in the Introduction, this shows up in the dependence of the weights on $a$ despite it corresponding to a horizon other than the target one. The single horizon approach resolves this issue since it removes dependence on horizons $h < H$ by construction.

Second, a shock receives greater weight if there are other shocks with impulse responses of similar shape. For instance, if the impulse responses of variable $i$ to shocks $2, \ldots, N$ are identical (or close) up to scale, then they will receive a large weight, potentially even if the impulse response to each of them individually is smaller than the impulse response to shock 1. The example in the Introduction showed that this can be a powerful force leading to counterintuitive results, with zero weight on shock 1 when $b = 0$ despite it being the only shock with a non-zero response at the target horizon. This is analogous to how the principal component of a set of variables depends both on the relative variance of the variables and their correlation structure. In the multiple horizon problem, identification thus depends on the similarity in impulse response shapes across all shocks even though the econometrician may not intend to impose such assumptions. This force is not present in the single horizon problem, since $\mathcal{H}$ is a singleton.

Using Lemma 1, we now characterize the single and multiple horizon problems. We will denote the solution to the single horizon problem with $\mathcal{H} = \{h\}$ by $\theta^S_h$ and the solution to the multiple horizon problem by $\theta^M_{\mathcal{H}}$. When $\mathcal{H} = \{0, \ldots, H\}$, we denote the solution by $\theta^M_{0:H}$.

**Single Horizon Problem.** In the single horizon case, we can tightly characterize the max-share solution in terms of the true shocks.

**Proposition 1.** *With $\mathcal{H} = \{h\}$ and $\widetilde{\Psi}_h = \Psi_h$, the solution to (6) is:*

$$\theta_h^S = \frac{\Psi_h' \delta_i}{\|\delta_i' \Psi_h\|} \tag{10}$$

*(up to a sign normalization).*

Proposition 1 implies that the weight on shocks $j$ in the single horizon case satisfies $\theta_{h,j}^S \propto \Psi_{h,ij}$. In other words, the max-share shock is a weighted average of the true shocks, where the weight on shock $j$ is proportionate to the size of impulse response of target variable $i$ to that shock at horizon $h$. Therefore, the solution to the single horizon problem will put a large weight on the true shock of interest if and only if the impulse response of variable $i$ to that shock is substantially larger than its response to all the other shocks combined. When identifying technology shocks or news shocks, this is typically the underlying assumption—the shock of interest is taken to be the only one to generate a persistent response in the target variable, hence generating a relatively large impulse response at some medium or long horizon.

Nevertheless, Proposition 1 makes clear that there are conditions under which the single horizon problem might fail to (approximately) identify the correct shock, as is the case in any identification scheme. In particular, the single horizon max-share problem identifies a shock that is necessarily contaminated by other shocks unless the impulse response of the target variable $i$ to those shocks is exactly zero at horizon $h$. While it is infeasible to check the relative size of the true impulse responses in empirical applications, the conditions present an object, i.e., the relative size of the impulse response to the target variable at a single horizon, that an econometrician might plausibly be able to form a prior over. One can also easily check if these conditions hold in a given model under reasonable parameterizations.

**Multiple Horizon Problem with Two Shocks.** For the multiple horizon problem, analytical results are generally unavailable. However, we provide intuition from the special case with two shocks, then provide simulation evidence for more general cases.

**Lemma 2.** *With $N = 2$ shocks, the ratio of the weights in the solution to (6) is:*

$$\frac{\theta_{\mathcal{H},1}^M}{\theta_{\mathcal{H},2}^M} = \frac{\vartheta + \sqrt{\vartheta^2 + 4}}{2}, \tag{11}$$

*where*

$$\vartheta \equiv \frac{\|\psi_{\mathcal{H},1}\|^2 - \|\psi_{\mathcal{H},2}\|^2}{\psi_{\mathcal{H},1} \cdot \psi_{\mathcal{H},2}} = \frac{1}{\cos\alpha_{12}} \left( \frac{\|\psi_{\mathcal{H},1}\|}{\|\psi_{\mathcal{H},2}\|} - \frac{\|\psi_{\mathcal{H},2}\|}{\|\psi_{\mathcal{H},1}\|} \right) \tag{12}$$

*and $\psi_{\mathcal{H},j}$ is a vector whose kth entry is the impulse response of variable $i$ to shock $j$ at the kth horizon, $h_k$, of $\mathcal{H}$, as before.*

The solution (11) is increasing in $\vartheta$, which captures the difference in the contributions of each shock to the FEV of variable $i$, normalized by a dot product term. The numerator of $\vartheta$ indicates that shock 1 receives more weight if it generates a relatively larger impulse response overall across horizons $h \in \mathcal{H}$. The denominator implies that the weights are more even if the shape of the two impulse responses is closer. In the limit, with $\psi_{\mathcal{H},1}$ orthogonal to $\psi_{\mathcal{H},2}$, i.e., $\cos\alpha_{12} = 0$, all the weight is placed on the shock with the larger impulse response. The variable $\vartheta$ thus captures the role of the magnitude and correlation across the impulse response of variable $i$ to each shock. If shock 1 is the true shock, then one should choose $\mathcal{H}$ to maximize the relative magnitude of the impulse response to shock 1 ($\|\psi_{\mathcal{H},1}\| / \|\psi_{\mathcal{H},2}\| - \|\psi_{\mathcal{H},2}\| / \|\psi_{\mathcal{H},1}\|$) and minimize the similarity between the impulse responses' shapes ($\cos\alpha_{12}$) in order to maximize the weight on shock 1 (assuming $\|\psi_{\mathcal{H},1}\| > \|\psi_{\mathcal{H},2}\|$).

The solution to the single horizon problem can also be expressed as (11), but with $\vartheta = \|\psi_{h,1}\| / \|\psi_{h,2}\| - \|\psi_{h,2}\| / \|\psi_{h,1}\|$. This is equivalent to (12) with $\mathcal{H} = \{h\}$ and $\cos\alpha_{12} = 1$. In other words, the multiple horizon approach raises two concerns relative to the single horizon one. Not only could the relative size of the impulse response to the true shock turn out to be smaller when taken over the set of horizons $\mathcal{H}$, but the similarity between the impulse responses, as captured by $\cos\alpha_{12}$, could also alter the weight on the true shock.

More concretely, consider a model with stationary demand shocks and permanent supply shocks, which we will study more closely in Section 3. The above results show that increased persistence in the demand shock weakens identification through three channels in the multiple horizon approach with $\mathcal{H} = \{0, \ldots, H\}$. First, it increases the relative size of the demand shock at the target horizon. Second, it also increases the relative size of the demand shock at all prior horizons. Third, it makes the shape of the impulse response to the demand shock closer to that corresponding to the supply shock. With the single horizon approach, the latter two forces are no longer in effect.

**Corollary 1.** *Consider the case with $N = 2$ shocks. Suppose the impulse responses are positive, i.e., $\Psi_{h,ij} > 0$ for all $h \in \mathcal{H}$ and $j \in \{1, 2\}$, and $\psi_{\mathcal{H},1} \cdot \psi_{\mathcal{H},2} \neq 0$. Then the multiple*

*horizon weights must be bounded by the single horizon weights:*

$$\theta^M_{\mathcal{H},j} \leq \max_{h\in\mathcal{H}} \theta^S_{h,j} \tag{13}$$

*for* $j \in \{1, 2\}$.

Corollary 1 provides sufficient conditions for the multiple horizon weights to be bounded by the corresponding single horizon ones. The condition that the impulse responses are positive simply requires that the impulse response function does not cross zero at some horizon, which is true, for instance, when distinguishing news and surprise TFP shocks or demand and supply shocks in many models. As we will see below, the condition is not necessary, and the bound, in fact, holds for a wide range of impulse responses.

**Simulation Evidence for Multiple Horizon Problem.** In the absence of general theoretical results for the multiple horizon problem with $N > 2$, we now provide simulation evidence to show that the insights from the Corollary 1 hold more generically.

Specifically, we consider impulse responses for the target variable $i$ of the form:

$$\Psi_{h,ij} = \mu_j + (-1)^{\varsigma}\rho_j^h + \epsilon_{h,j} \tag{14}$$

where $\mu_j \sim \text{Uniform}(-1, 1)$, $\rho_j \sim \text{Uniform}(0, 1)$, $\varsigma \sim \text{Bernoulli}(0.5)$, and $\epsilon_{h,j} \sim \mathcal{N}(0, 0.1^2)$. These impulse responses have shape resembling that of an AR(1) process with persistence $\rho_j$, but deviate at each horizon by $\epsilon_{h,j}$ and converge to some long-run response $\mu_j$. This allows for arbitrary signs and relative sizes of impulse responses, while keeping them empirically plausible.[3] We generate $10^6$ draws and, for each draw, check if (13) holds for shock $j = 1$.

Figure 2 shows that for all values of $N$ and $H$ considered, the single horizon problem provides bounds for the multiple horizon weights for at least 92 percent of the draws. The fraction is increasing in both the number of shocks, $N$, and the horizon, $H$. It approaches 1 for large $H$ and exceeds 0.99 for empirically relevant choices of $N$ and $H$. These probabilities increase even further if we restrict the impulse responses to be positive as in Corollary 1 (taking $\mu_j \sim \text{Uniform}(0, 1)$, $\varsigma = 0$, and $\epsilon_{h,j} = 0$), with $P\left[\theta^M_{0:H,1} \leq \max_{h\in\{0,\ldots,H\}} \theta^S_{h,1}\right] > 0.994$ for all $N$ and $H$ we consider.[4]

For practitioners, the above results suggest that one is likely do at least as well using the single horizon problem with appropriately chosen $h$ rather than the multiple horizon problem with $\mathcal{H} = \{0, \ldots, H\}$. In many cases, such as our prototypical example in Figure 1,

[3] We obtain similar results taking $\epsilon_{h,j} = 0$ or $\Psi_{h,ij} \sim \text{Uniform}(-1, 1)$.

[4] Furthermore, we find $\left|\theta^M_{0:H,1}/\theta^M_{0:H,2}\right| \leq \max_{h\in\{0,\ldots,H\}} \left|\theta^S_{h,1}/\theta^S_{h,2}\right|$ for all draws.

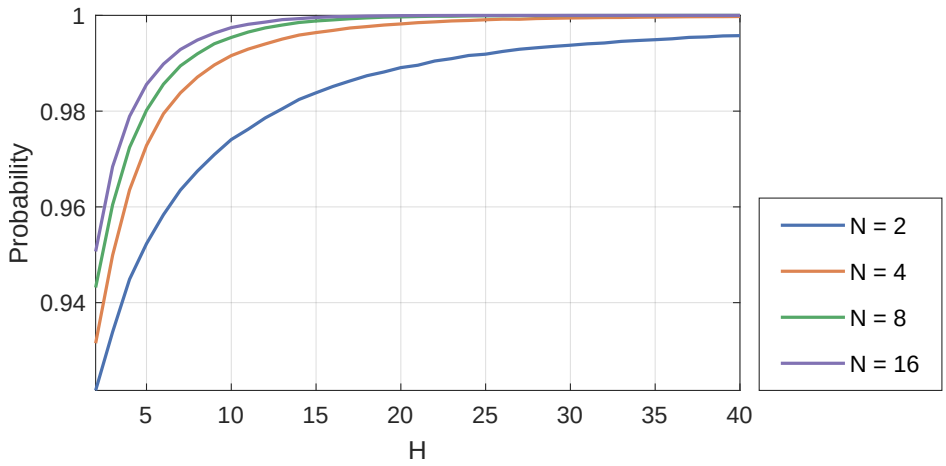


Figure 2: Probability that $\theta^M_{0:H,1} \leq \max_{h\in\{0,\ldots,H\}} \theta^S_{h,1}$ from simulated impulse responses.

the choice of $h$ supported by theory is simply $h = H$. Even though estimation uncertainty, which we have abstracted from for clarity, could lead to a preference for a different choice of $h$ or even a set of $h$'s, it is unlikely that including all horizons from 0 to $H$, as is standard in the literature, would be optimal.

# 3 Illustrative Example: Supply and Demand

We now illustrate the results from Section 2 in the context of a simple demand and supply system that can be expressed as a bivariate VAR(1). We assume the reduced form parameters are known and identify the supply shock using max-share identification on output for a large but finite $H$. The approach follows Francis et al. (2014), who target labor productivity instead of output.

## 3.1 Setup

Consider the simple demand and supply system:

$$q^s_t = \gamma^s p_t + \eta^s_t \tag{15}$$

$$q^d_t = -\gamma^d p_t + \eta^d_t \tag{16}$$

where $q^s_t$, $q^d_t$, and $p_t$ denote the quantity demanded, quantity supplied, and price, respectively, in logs. The supply shock $\eta^s_t$ follows a random walk and demand shock $\eta^d_t$ follows a stationary AR(1) process.

Using the fact that $q_t = q^s_t = q^d_t$ in equilibrium and normalizing coefficient on the supply

shock innovations, we can write the system as a VAR(1):

$$Y_t = GRG^{-1}Y_{t-1} + GS\varepsilon_t, \tag{17}$$

where

$$Y_t = \begin{bmatrix} q_t \\ p_t \end{bmatrix}, R = \begin{bmatrix} 1 & 0 \\ 0 & \rho^d \end{bmatrix}, S = \begin{bmatrix} 1 & 0 \\ 0 & \sigma^d \end{bmatrix}, \text{ and } G = \frac{1}{\gamma^s + \gamma^d} \begin{bmatrix} \gamma^d & \gamma^s \\ -1 & 1 \end{bmatrix}.$$

The innovations $\varepsilon_t$ are iid standard normal and $\rho^d$ is the persistence of the demand shock.

Since the supply shock is permanent but the demand shock is transitory, the supply shock dominates the FEV of $q_t$ at the infinite horizon. Therefore, with known reduced form parameters, the long-run identification of Blanchard and Quah (1989), i.e., setting $\mathcal{H} = \{0, \ldots, H\}$ with $H \to \infty$, correctly identifies the supply shock. Francis et al. (2014) consider a finite $H$ in order to account for uncertainty in the reduced form parameters and contamination from low frequency trends.

## 3.2 Analytics

The true impulse response at horizon $h$ to a one unit innovation in the $j$th shock is:

$$\Psi_h = GR^hS = \frac{1}{\gamma^s + \gamma^d} \begin{bmatrix} \gamma^d & \sigma^d(\rho^d)^h\gamma^s \\ -1 & \sigma^d(\rho^d)^h \end{bmatrix}. \tag{18}$$

The single horizon problem with $\mathcal{H} = \{h\}$ obtains the impulse response to a combination of supply and demand shocks:

$$\theta_h^S \propto \begin{bmatrix} \gamma^d \\ \sigma^d(\rho^d)^h\gamma^s \end{bmatrix}. \tag{19}$$

This is simply the transpose of the top row of the impulse response matrix (18), as shown in Proposition 1.

The demand innovation, $\varepsilon_t^d$ receives a heavier weight if it is more persistent ($\rho^d$) and volatile ($\sigma^d$) since these increase response of the demand shock, $\eta_t^d$, at horizon $h$ periods after the initial innovation. In addition, the demand shock receives a heavier weight if it is relatively less elastic ($\gamma^s/\gamma^d$). If the price elasticity of demand, $\gamma^d$, is smaller, then quantities, $q_t$, respond more to the demand shock, $\eta_t^d$. In this simple example, the horizon $h$ only shows up as an exponent on $\rho^d$. As $h \to \infty$, all the weight is placed on the supply shock. However,

the relative weights for finite $h$ are parameter-dependent.

We can contrast the single horizon approach with the multiple horizon problem, which yields the solution (11) with:

$$\vartheta = \frac{\gamma^s \sigma^d}{\gamma^d} \frac{1 + (\rho^d)^{H+1}}{1 + \rho^d} - (H+1) \frac{\gamma^d}{\gamma^s \sigma^d} \frac{1 - \rho^d}{1 - (\rho^d)^{H+1}}. \tag{20}$$

For the supply shock to be correctly identified, we require the ratio of weights on the demand and supply shocks to converge to zero, i.e., $\theta^M_{0:H,d}/\theta^M_{0:H,s} \to 0$. This is indeed the case as $H \to \infty$. However, there is no guarantee that the weight on the supply shock will be large with finite $H$. For example, with $\rho^d = 0$, the multiple horizon problem will only place majority of the weight on the supply shock ($\theta^M_{0:H,d}/\theta^M_{0:H,s} < 1$) for $H > \left(\frac{\gamma^s \sigma^d}{\gamma^d}\right)^2 - 1$. The ratio $\frac{\gamma^s \sigma^d}{\gamma^d}$ captures the relative size of the demand shock impulse response on impact, highlighting the persistent influence of short horizons on the multiple horizon weights even when the demand shock itself displays no persistence. In contrast, the single horizon problem places all the weight on the true supply shock for any target horizon $h > 0$.

## 3.3 Numerical Example

As a numerical proof-of-concept, we set:

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

and assume as before that the supply shock follows a random walk with unit variance. We maximize the contribution to output at horizon $H = 40$ to identify the supply shock. We compare the typical multiple horizon implementation in the literature with $\mathcal{H} = \{0, \ldots, H\}$ to the single horizon case with $\mathcal{H} = \{H\}$. As with the discussion above, we take the reduced form parameters as given.

Figure 3 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 3 shows that the true supply shock only accounts for less than a third of the FEV, an indication that multiple horizon max-share is unlikely to perform well.

The multiple horizon 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

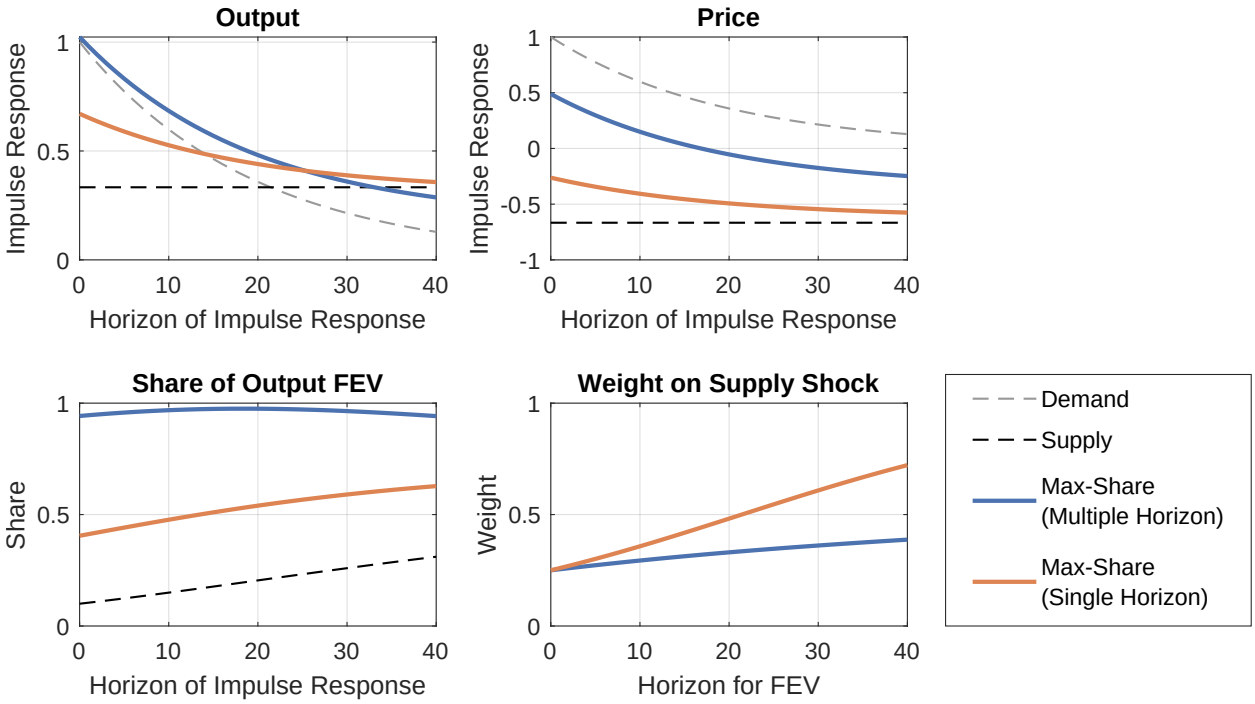


Figure 3: Results from numerical demand and supply example. **Top row:** Impulse responses of output and price; **Bottom left:** Contribution of shocks to FEV of output; **Bottom right:** Weight on supply shock with max-share identification for different target horizons. Dashed lines correspond to true shocks and solid lines correspond to identified max-share shocks.

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 $H$, the improvement is relatively slow, with the weight at $H = 0$ already at 0.25.

The single horizon approach brings the max-share shock closer to the true supply shock, with a weight of 0.72. Nevertheless, the behavior of the single horizon max-share and true supply shocks still differ. The response of price to the single-horizon max-share shock is negative but relatively small. In particular, the impulse response on impact implies an elasticity of demand, $\gamma^d$, of 2.6, over five times the size of its true value of 0.5. The contribution to the output FEV remains more than double the true share across horizons 0 to 40.

The weight on the supply shock increases much faster with $H$ in the single horizon case since the solution is not contaminated by the large response of output to demand shocks at short horizons. When we set $H = 80$, as the weight on the true supply shock rises to 0.95 for the single horizon problem but remains at 0.48 for the multiple horizon problem.

Overall, the above example emphasizes how the impulse responses at short horizons can

have a substantial impact on the FEV even for relatively long horizons, thus impacting the multiple horizon max-share identification approach. The single horizon problem improves upon this by focusing narrowly on a single horizon for which the econometrician may have a stronger prior on the relative magnitudes of the true impulse responses. However, unless the true impulse response strongly dominates the remaining shocks at that horizon, one should remain cautious in making quantitative statements.

# 4 Empirical Application: News Shocks

To show how the above analysis matters empirically, we now consider a common application of max-share identification—estimating the effect and importance of TFP news shocks. This follows a large literature (Beaudry and Portier, 2006; Barsky and Sims, 2011; Schmitt-Grohé and Uribe, 2012) that seeks to identify shocks that predict future productivity without being related to current or past fundamentals. Since the effect of 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 identification that target TFP at a relatively long horizon to identify the news shock.

## 4.1 Data and Estimation

We follow Kurmann and Sims (2021) and estimate an 8-variable VAR with utilization-adjusted TFP from Fernald (2014), the S&P 500 index, real consumption per capita, real GDP per capita, real investment per capita, hours per capita, GDP deflator inflation, and the federal funds rate. 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.[5]

To identify the shock, we follow the approach of Kurmann and Sims (2021) and use the multiple horizon problem with $\mathcal{H} = \{0, \ldots, 40\}$, i.e., finding the shock that contributes the most to the FEV at horizon $H = 40$. We compare this benchmark to the single horizon problem that similarly takes $h = 40$.

[5]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. Our results are also robust to including data through 2023Q4 or estimating the smaller 4-variable VAR in Kurmann and Sims (2021), which includes only TFP, consumption, hours, and inflation.

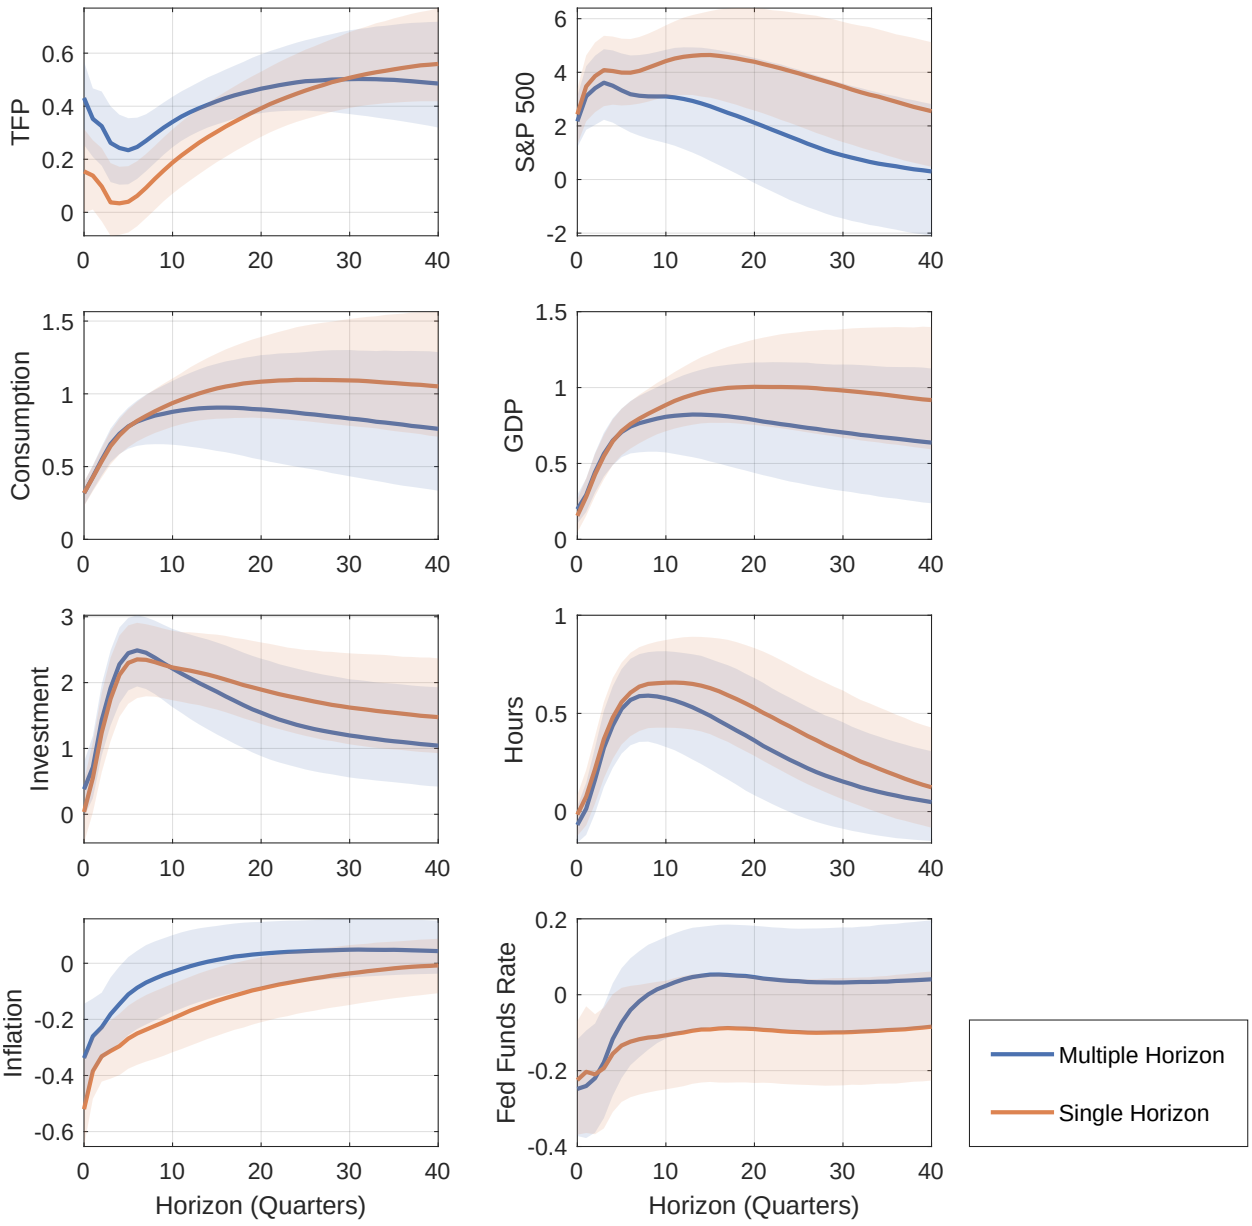


Figure 4: Posterior estimates for impulse responses to identified max-share shock. **Solid lines:** Median response; **Shaded regions:** 68% error bands.

## 4.2 Results

Figure 4 shows the estimated impulse responses and Figure 5 shows the contribution of the max-share shocks to the FEV of each variable. While the 68% error bands overlap, the point estimates are notably different.

The single horizon shock generates an impulse response of TFP that is smaller at short horizons but more persistent and larger at longer horizons than that of the multiple horizon shock. Similarly, while the multiple horizon shock accounts for around 1/3 of the FEV at

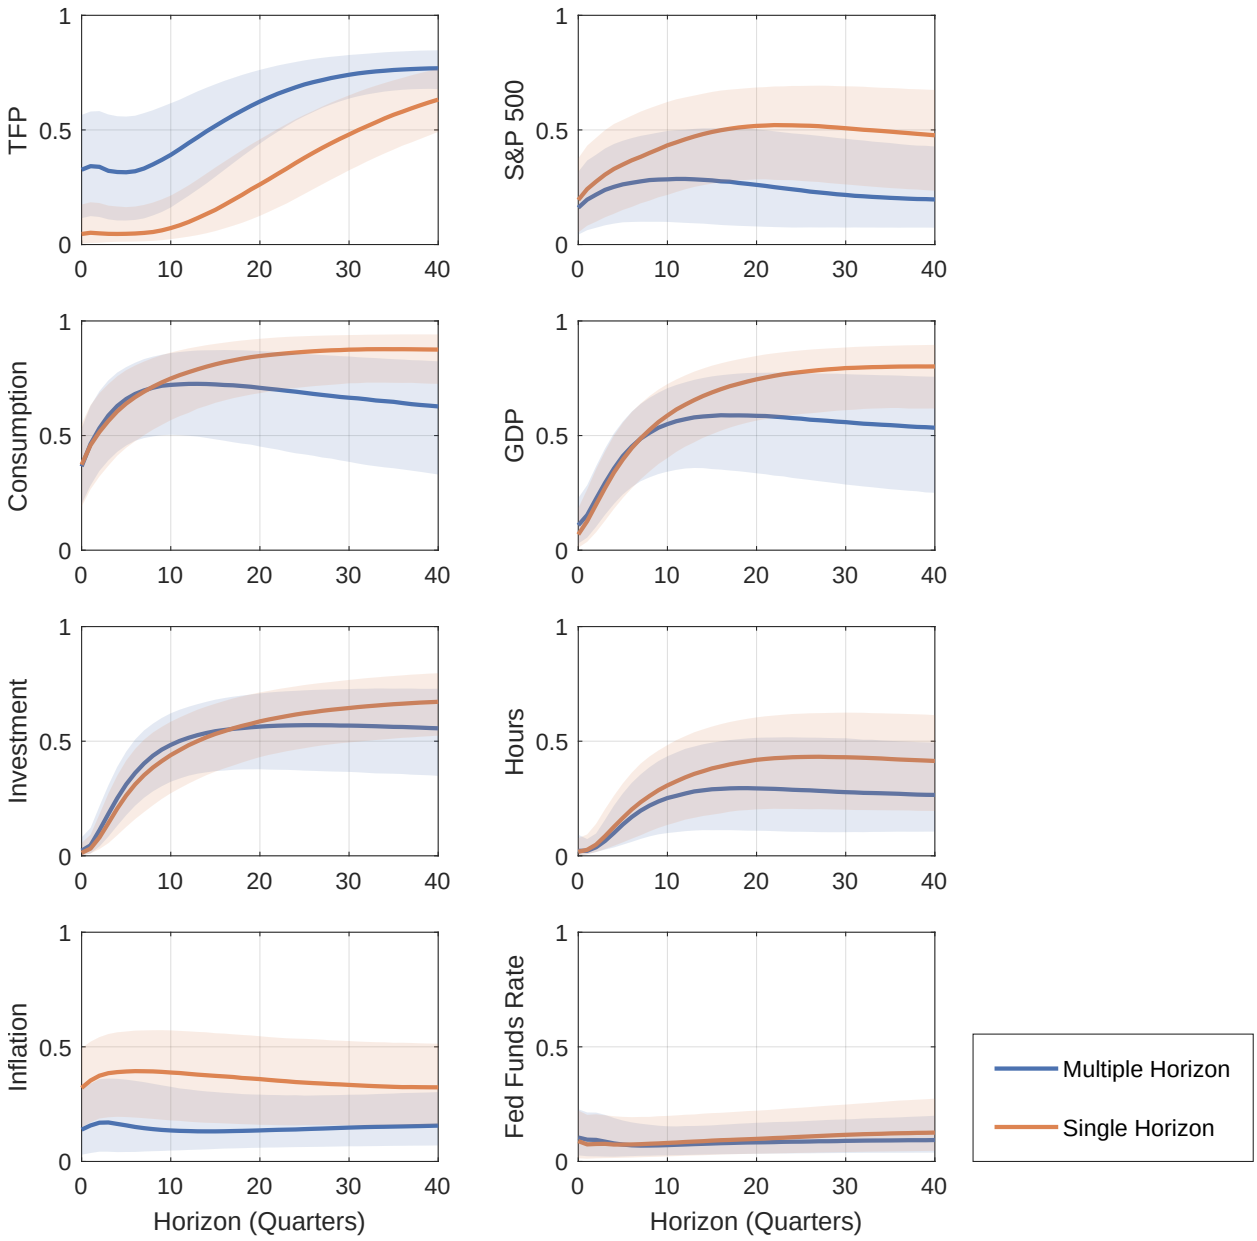


Figure 5: Posterior estimates for contribution of identified max-share shock to FEV. **Solid lines:** median; **Shaded regions:** 68% error bands.

short horizons, the single horizon shock accounts for less than 5% of the FEV. The gap narrows somewhat at longer horizons, with contributions of 77% and 63%, respectively, 40 quarters out. The small impulse response and contribution of the single horizon shocks at short horizons is more consistent the interpretation of news shocks having a delayed effect on TFP, which is the case in models such as those in Kurmann and Sims (2021) or Schmitt-Grohé and Uribe (2012) as well as empirical estimates from Beaudry and Portier (2006).

Even though the response of TFP through horizon $h = 30$ is smaller under single horizon

max-share identification, the responses of all the other variables are larger than with multiple horizon max-share identification. For example, inflation falls by over 1.5 times as much on impact and real GDP increases 1.5 times as much 40 quarters out.[6] Similarly the single horizon shock tends to contribute more to the FEV of most other variables, especially at longer horizons. For example, the single horizon shock accounts for 80% of the FEV of GDP growth as compared to 53% for the multiple horizon shock.

These results highlight the advantages and limitations of using single horizon max-share identification instead of the usual multiple horizon approach in the literature. On the one hand, using the single horizon approach reduces contamination of our identified shock by other shocks that have relatively transitory effects. For example, suppose there is measurement error, as emphasized by Kurmann and Sims (2021). Assuming the measurement error is transitory, one would expect its relative impact on TFP to be larger at short horizons without having any actual impact on other variables. Alternatively, Cascaldi-Garcia and Galvao (2021) argue that financial uncertainty shocks contaminate the max-share identified TFP news shocks, leading to a smaller impulse response of real activity to the identified shock. This is consistent with the differences between the single and multiple horizon identification results. Because the multiple horizon identification places a greater weight on transitory shocks such as measurement error or financial uncertainty, it generates a TFP response that is larger on impact and less persistent. In addition, other variables have smaller responses since these contaminating shocks generate responses of smaller magnitude or opposite signs.

On the other hand, single horizon identification remains influenced by any other shocks that continue to have an effect at horizon $h = 40$. A shock that induces an increase in investment could generate a persistent rise in TFP. Even though such a shock is likely not what one has in mind as an exogenous technological news shock in a typical model, it could generate a substantial response in TFP, thus causing it to contaminate the identified single horizon max-share shock. For instance, Justiniano, Primiceri, and Tambalotti (2011) and Ben Zeev and Khan (2015) find quantitatively important roles for marginal efficiency of investment shocks and investment-specific technology new shocks, respectively. Nevertheless, Proposition 1 makes clear that the amount of contamination depends precisely on the relative size of the true impulse responses at the target horizon, making the underlying identification assumptions transparent.

Figure 6 compares our baseline results with $H = 40$ to those with $H = 80$, further highlighting the difference between the single and multiple horizon approaches. The impulse

[6]The relative size of the single horizon responses is consistent with estimates from Miranda-Agrippino et al. (2022), who also estimate a structural VAR, but identify news shocks using patent applications as an external instrument.

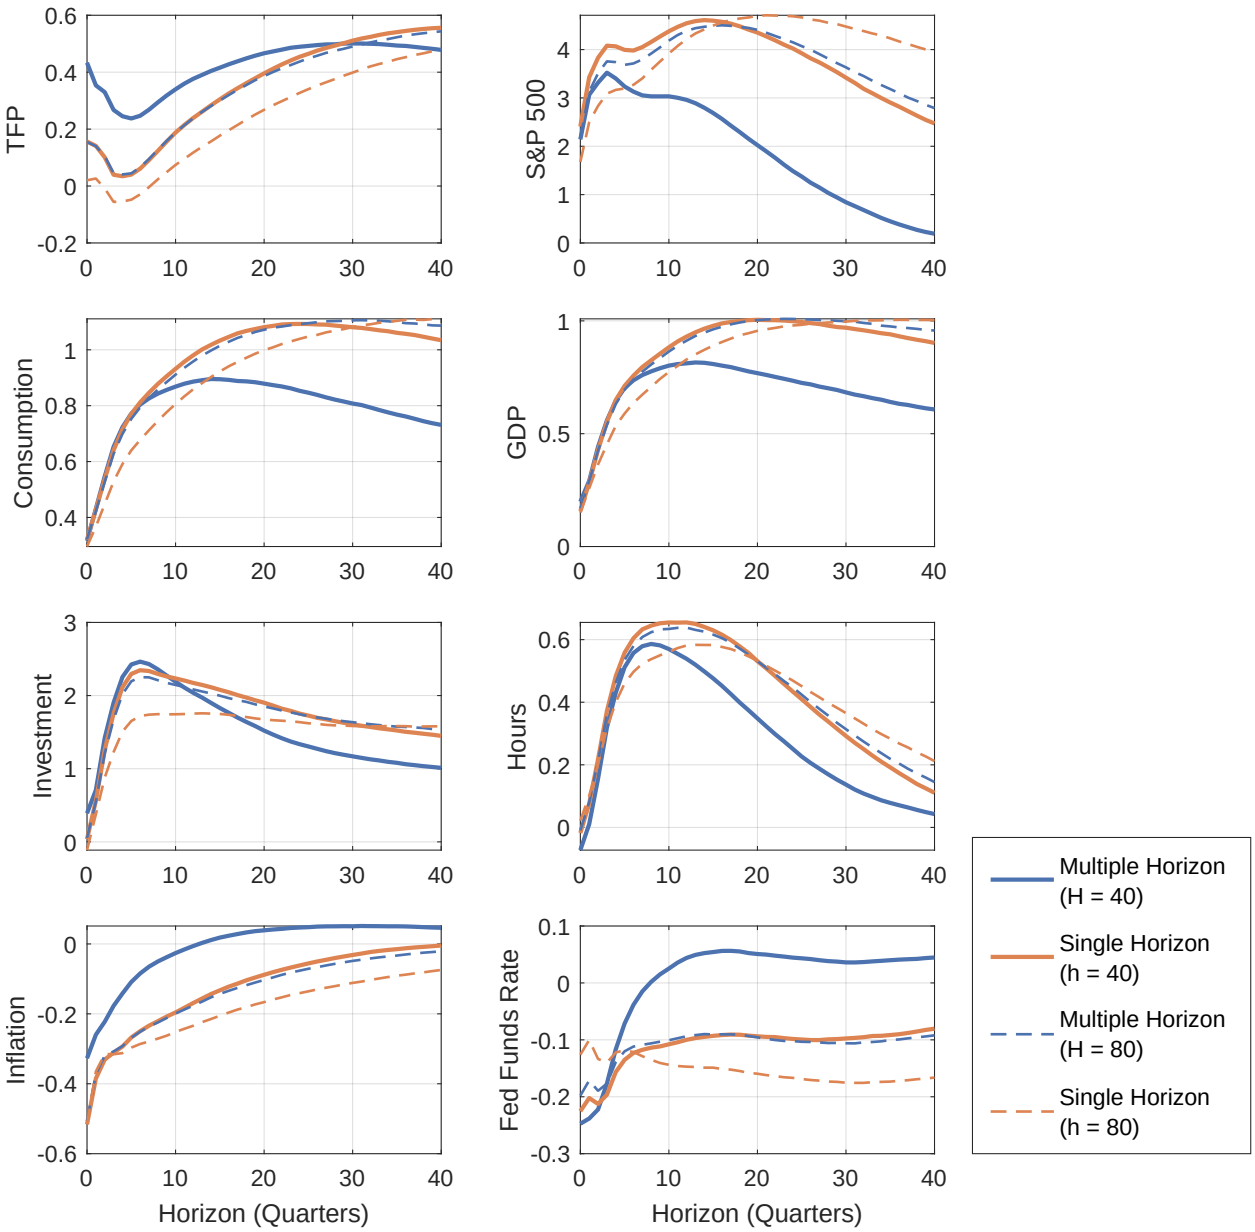


Figure 6: Posterior median for impulse responses to identified max-share shock for different horizons. **Solid lines:** $H = 40$; **Dashed lines:** $H = 80$.

responses from the two approaches become closer with $H = 80$. These impulse responses in turn are relatively similar to the single horizon impulse response for $H = 40$. In other words, the single horizon approach appears more robust to the choice of horizon. Intuitively, the contamination from short horizons in the multiple horizon approach decreases as we consider a longer horizon, since those horizons constitute a smaller part of the summation (7). For a practitioner, this provides another argument for using the single horizon approach. It continues to serve purpose described by Francis et al. (2014) of standing in for long-run

restrictions while avoiding the potential drawbacks of taking $H \to \infty$. However, the results are less dependent on the exact choice of $H$, thus allowing one to choose a smaller $H$ and better guarding against the issues from the long-horizon impulse response estimates.

# 5 Conclusion

Despite the increasing popularity of max-share identification, the interpretation of the resulting identified shock is less clear than one might expect. In particular, even though max-share identification, as commonly implemented, targets the FEV at a given horizon $H$, it depends crucially on the impulse responses at all horizons $h \leq H$. Therefore, it inadvertently conflates the relative size of impulse responses to different shocks across all these horizons even if one is only willing to place assumptions on the responses at horizon $H$. Moreover, the weights depend on the shapes of the impulse responses relative to each other, further obscuring the interpretation.

We propose as an alternative a max-share identification that focuses only on one horizon and thereby avoids these pitfalls, unlike the multiple horizon problem, which sums across all horizons $h \in \mathcal{H}$ when constructing the FEV. When the data-generating process is a structural VAR, the resulting weights on the true shocks are exactly proportional to the relative size of the impulse response of the target variable to each of the shocks, allowing one to easily discern conditions under which the true shock will be well-approximated by the identified shock. If, at horizon $h$, the target variable still has substantial responses to multiple shocks, then the our results reveal the exact degree by which the identified shock is contaminated.

These insights are illustrated in an empirical news shock application. On the one hand, the single horizon problem produces a TFP impulse response that is closer to theory, suggesting that some measurement error or other auxiliary shocks might have been downweighted relative to the multiple horizon approach. On the other hand, it further increases the contribution to real activity, potentially due to a greater weight not only on TFP news shocks, but also other shocks that have an effect on TFP at the 10-year horizon.

# References


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 (2023). Risky Business Cycles. Working paper.

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.

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

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.

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

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.

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.

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

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2011). Investment Shocks and the Relative Price of Investment. *Review of Economic Dynamics 14*(1), 102–121.

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

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.

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., S. H. Hoke, and K. Bluwstein (2022). Patents, News, and Business Cycles. Working paper.

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.

Ramey, V. A. (2016). Macroeconomic Shocks and Their Propagation. *Handbook of Macroeconomics 2*, 71–162.

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

Tao, T. (2012). *Topics in Random Matrix Theory*, Volume 132. American Mathematical Society.

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.

# A Proofs

## A.1 Lemma 1

*Proof.* The lemma follows from:

$$\delta_i' \left[ \sum_{h \in \mathcal{H}} \widetilde{\Psi}_h \theta \theta' \widetilde{\Psi}_h' \right] \delta_i = \sum_{h \in \mathcal{H}} \mathrm{Tr} \left( \theta' \widetilde{\Psi}_h' \delta_i \delta_i' \widetilde{\Psi}_h \theta \right) = \theta' \left[ \sum_{h \in \mathcal{H}} \widetilde{\Psi}_h' \delta_i \delta_i' \widetilde{\Psi}_h \right] \theta. \tag{21}$$

Note that the objective is a scalar and thus equals its own trace. Because the trace is a linear mapping, the objective becomes $\sum_{h \in \mathcal{H}} \mathrm{Tr} \left( \delta_1' \widetilde{\Psi}_h \theta \theta' \widetilde{\Psi}_h' \delta_1 \right)$. This, together with the fact that $\mathrm{Tr} \left( \delta_1' \widetilde{\Psi}_h \theta \theta' \widetilde{\Psi}_h' \delta_1 \right) = \mathrm{Tr} \left( \theta' \widetilde{\Psi}_h' \delta_1 \delta_1' \widetilde{\Psi}_h \theta \right)$ for each $h$, leads to the first equality. The second equality holds again by the linearity of the trace operator. □

## A.2 Proposition 1

*Proof.* Denote $\Psi_h' \delta_i \delta_i' \Psi_h$ by $\Omega_h$. First, note that $\Omega_h$ is always of rank one since each row is a scale multiple of its $i$th row $\delta_i' \Psi_h$. It thus has only one non-zero eigenvalue, which is obtained by the following fact,

$$\Omega_h \theta_h^S = \frac{\Psi_h' \delta_i \delta_i' \Psi_h \Psi_h' \delta_i}{\| \delta_i' \Psi_h \|} = \| \delta_i' \Psi_h \|^2 \theta_h^S. \tag{22}$$

Because $\| \delta_i' \Psi_h \|^2 > 0$ for nontrivial $\delta_i' \Psi_h$, $\theta_h^S$ solves the eigenproblem for (7) with $\mathcal{H} = \{h\}$. Therefore, the solution to (6) is $\theta_h^S$ (up to a sign normalization) by Lemma 1. □

## A.3 Lemma 2

*Proof.* With two shocks, the solution for $\theta$ corresponds to the eigenvector associated with

$$\Phi_{\mathcal{H}} \equiv \sum_{h \in \mathcal{H}} \Psi_h' \delta_i \delta_i' \Psi_h = \begin{bmatrix} \phi_{11} & \phi_{12} \\ \phi_{12} & \phi_{22} \end{bmatrix} \text{ where } \phi_{jj'} = \sum_{h \in \mathcal{H}} \Psi_{h,ij} \Psi_{h,ij'}. \tag{23}$$

The eigenvector $\theta^*$ satisfies $(\Phi_{\mathcal{H}} - \lambda I)\theta^* = 0$, which implies:

$$\lambda = \phi_{11} + \phi_{12} \frac{\theta_2^*}{\theta_1^*} = \phi_{22} + \phi_{12} \frac{\theta_1^*}{\theta_2^*}. \tag{24}$$

Multiplying throughout by $\theta_2^*/\theta_1^*$ and dividing by $\phi_{12}$, we have:

$$\left(\frac{\theta_1^*}{\theta_2^*}\right)^2 - \vartheta\left(\frac{\theta_1^*}{\theta_2^*}\right) - 1 = 0 \tag{25}$$

The quadratic has two solutions. Substituting these back into (24), we can verify that the larger eigenvalue corresponds to:

$$\frac{\vartheta + \sqrt{\vartheta^2 + 4}}{2}. \tag{26}$$

□

## A.4 Corollary 1

*Proof.* Without loss of generality, consider $j = 1$ and suppose that

$$\arg\max_{h \in \mathcal{H}} |\theta_{h,1}^S| = H. \tag{27}$$

Since $\Psi_{h,ij} \geq 0$, we can normalize $\theta_{h,j}^S$ to be positive for all $h \in \mathcal{H}$.

Consider an alternative set of impulse response functions $\widehat{\Psi}_h$ satisfying:

$$\widehat{\Psi}_{h,i1}\widehat{\Psi}_{h,i2} = \Psi_{h,i1}\Psi_{h,i2} \tag{28}$$

$$\frac{\widehat{\Psi}_{h,i1}}{\widehat{\Psi}_{h,i2}} = \frac{\Psi_{H,i1}}{\Psi_{H,i2}} \tag{29}$$

for all $h \in \mathcal{H}$. Solving the above equations yields:

$$\widehat{\Psi}_{h,i1}^2 = \frac{\Psi_{h,i1}\Psi_{h,i2}}{\widehat{\Psi}_{h,i2}}\widehat{\Psi}_{h,i1} = \frac{\Psi_{H,i1}}{\Psi_{H,i2}}\Psi_{h,i1}\Psi_{h,i2} > \frac{\Psi_{h,i1}}{\Psi_{h,i2}}\Psi_{h,i1}\Psi_{h,i2} = \Psi_{h,i1}^2 \tag{30}$$

and similarly $\widehat{\Psi}_{h,i2}^2 < \Psi_{h,i2}^2$. For the alternative impulse responses $\widehat{\Psi}_h$, because $\widehat{\theta}_{\mathcal{H}}^M = \widehat{\theta}_h^S$ by construction, the solutions to the multiple horizon problem and single horizon problem for any $h \in \mathcal{H}$ coincide by Lemma 3 (the linear span collapses to $\{\widehat{\theta}_h^S\}$), i.e., $\widehat{\theta}_{\mathcal{H}}^M = \widehat{\theta}_h^S = \theta_H^S$.

Comparing the expressions for $\vartheta$ corresponding to $\Psi_h$ and $\widehat{\Psi}_h$, we have:

$$\vartheta = \frac{\|\psi_{\mathcal{H},1}\|^2 - \|\psi_{\mathcal{H},2}\|^2}{\psi_{\mathcal{H},1} \cdot \psi_{\mathcal{H},2}} < \frac{\left\|\widehat{\psi}_{\mathcal{H},1}\right\|^2 - \left\|\widehat{\psi}_{\mathcal{H},2}\right\|^2}{\widehat{\psi}_{\mathcal{H},1} \cdot \widehat{\psi}_{\mathcal{H},2}} = \widehat{\vartheta}. \tag{31}$$

To see this, notice that the denominators are equal by construction. In addition,

$$\|\psi_{\mathcal{H},1}\|^2 = \sum_{h\in\mathcal{H}} \Psi_{h,i1}^2 < \sum_{h\in\mathcal{H}} \widehat{\Psi}_{h,i1}^2 = \left\|\widehat{\psi}_{\mathcal{H},1}\right\|^2 \tag{32}$$

and similarly $\|\psi_{\mathcal{H},2}\|^2 > \left\|\widehat{\psi}_{\mathcal{H},2}\right\|^2$. But we know that $\theta_{\mathcal{H},1}^M$ is increasing in $\vartheta$ since $\theta_{\mathcal{H},1}^M/\theta_{\mathcal{H},2}^M$ is increasing in $\vartheta$ from Lemma 2. Therefore, $\theta_{\mathcal{H},1}^M < \widehat{\theta}_{\mathcal{H},1}^M = \theta_{H,1}^S$. □

**Lemma 3.** *The solution to* (6) *lies in* $\text{span}\{\theta_h^S | h \in \mathcal{H}\}$*, where* $\theta_h^S$ *is defined as in* (10)*.*

*Proof.* Denote the cardinality of $\mathcal{H}$ by $|\mathcal{H}|$ and label the set of horizons $\mathcal{H} = \{h_0, \ldots, h_{|\mathcal{H}|}\}$. We can write:

$$\sum_{h\in\mathcal{H}} \Psi_h' \delta_i \delta_i' \Psi_h = \Phi_{\mathcal{H}} \Phi_{\mathcal{H}}',$$

where $\Phi_{\mathcal{H}}$ is a $N \times |\mathcal{H}|$ matrix whose $(j,k)$ entry is the response of variable $i$ to shock $j$ at horizon $h_k$, collecting the impulse responses from different shocks to the variable $i$ at different horizons as columns, in the same ordering as the horizon indices in $\mathcal{H}$.

By Lemma 1, we know that $\theta_{\mathcal{H}}^M$ is the eigenvector associated with the largest eigenvalue of $\Phi_{\mathcal{H}}\Phi_{\mathcal{H}}'$, $\lambda_1(\Phi_{\mathcal{H}}\Phi_{\mathcal{H}}')$, and satisfy the following equality

$$\left[\sum_{h\in\mathcal{H}} \Psi_h' \delta_i \delta_i' \Psi_h\right] \theta_{\mathcal{H}}^M = \lambda_1(\Phi_{\mathcal{H}}\Phi_{\mathcal{H}}')\, \theta_{\mathcal{H}}^M. \tag{33}$$

A rearrangement of the left-hand side of (33) yields

$$\left[\sum_{h\in\mathcal{H}} \Psi_h' \delta_i \delta_i' \Psi_h\right] \theta_{\mathcal{H}}^M = \sum_{h\in\mathcal{H}} w_h \Psi_h' \delta_i, \tag{34}$$

where $w_h = (\delta_i'\Psi_h)\,\theta_{\mathcal{H}}^M$ is a $h$-specific scalar weight.

On the other hand, $\lambda_1(\Phi_{\mathcal{H}}\Phi_{\mathcal{H}}') > 0$ as long as there exists a $h_1$ such that $\|\delta_i'\Psi_{h_1}\| > 0$. To see this, consider another $h_2 \in \mathcal{H}$, by the classical (dual) Weyl inequality in e.g., Exercise 1.3.5. in Tao (2012), we have $\lambda_1\left(\Psi_{h_1}'\delta_i\delta_i'\Psi_{h_1} + \Psi_{h_2}'\delta_i\delta_i'\Psi_{h_2}\right) \geq \lambda_1\left(\Psi_{h_1}'\delta_i\delta_i'\Psi_{h_1}\right) + \lambda_N\left(\Psi_{h_1}'\delta_i\delta_i'\Psi_{h_2}\right) \geq \|\delta_i'\Psi_{h_1}\|^2 > 0$. Application of this argument recursively over $h \in \mathcal{H}$ leads to a strictly positive $\lambda_1(\Phi_{\mathcal{H}}\Phi_{\mathcal{H}}')$. This observation, along with (33) and (34) leads to an expression of $\theta_{\mathcal{H}}^M$ as a linear combination of $\theta_h^S$'s,

$$\theta_{\mathcal{H}}^M = \sum_{h\in\mathcal{H}} \frac{w_h \|\delta_i'\Psi_h\|}{\lambda_1(\Phi_{\mathcal{H}}\Phi_{\mathcal{H}}')} \theta_h^S.$$

□