



RESERVE BANK OF AUSTRALIA

# Shock-percentile Restrictions for SVARs

By Matthew Read

Research Discussion Paper 2026-01

The Discussion Paper series is intended to make the results of the current economic research within the Reserve Bank of Australia (RBA) available to other economists. Its aim is to present preliminary results of research so as to encourage discussion and comment. Views expressed in this paper are those of the authors and not necessarily those of the RBA. However, the RBA owns the copyright in this paper.

© Reserve Bank of Australia 2026

Apart from any use permitted under the *Copyright Act 1968*, and the permissions explicitly granted below, all other rights are reserved in all materials contained in this paper.

All materials contained in this paper, with the exception of any Excluded Material as defined on the RBA website, are provided under a Creative Commons Attribution 4.0 International License. The materials covered by this licence may be used, reproduced, published, communicated to the public and adapted provided that there is attribution to the authors in a way that makes clear that the paper is the work of the authors and the views in the paper are those of the authors and not the RBA.

For the full copyright and disclaimer provisions which apply to this paper, including those provisions which relate to Excluded Material, see the RBA website.

## **Enquiries**

Phone : +612 9551 8111

Email : [rbainfo@rba.gov.au](mailto:rbainfo@rba.gov.au)

Website : <https://www.rba.gov.au>

Some figures in this publication were generated using Mathematica.

ISSN 1448-5109 (Online)

# Shock-percentile Restrictions for SVARs

Matthew Read

Research Discussion Paper  
2026-01

March 2026

Economic Research Department  
Reserve Bank of Australia

For helpful comments, I thank Efram Castelnuovo, Jarkko Jääskelä, Gabriela Nodari, participants at the 2025 Workshop of the Australasian Macroeconomics Society, the 2025 OzMac Workshop and the Inaugural AE<sup>2</sup> Conference, and seminar participants at the University of Sydney and the Reserve Bank of Australia. I am particularly grateful to James Morley for discussing the paper at the Inaugural AE<sup>2</sup> Conference. The views expressed in this paper are those of the author and should not be attributed to the Reserve Bank of Australia. Any errors are the sole responsibility of the author.

Author: readm at domain rba.gov.au

External Communications: rbainfo@rba.gov.au

<https://doi.org/10.47688/rdp2026-01>

## **Abstract**

I propose identifying structural vector autoregressions using 'shock-percentile' restrictions. These restrictions require the realisation of a structural shock in a selected episode to lie in the tail of the shock's historical distribution, representing the belief that a relatively large shock has occurred. I argue that shock-percentile restrictions are an attractive alternative to imposing numeric bounds on shock magnitudes, which are difficult to credibly elicit. Simulations demonstrate the potential for shock-percentile restrictions to provide identifying information. In two empirical applications, I exploit shock-percentile restrictions to disentangle the relationship between uncertainty and real activity, and to sharpen identification of the macroeconomic effects of US monetary policy.

JEL Classification Numbers: C32, D80, E32, E44, E52

Keywords: monetary policy, narrative restrictions, set identification, structural vector autoregression, uncertainty

## Table of Contents

1.	Introduction	1
2.	Framework	3
2.1	SVAR	3
2.2	Identified sets	4
2.3	Shock-magnitude restrictions	5
3.	Challenges with Eliciting Bounds on Shocks	5
3.1	The RWZ algorithm	6
3.2	The bound elicitation procedure	6
3.3	Can the procedure be motivated from a Bayesian perspective?	9
4.	Shock-percentile Restrictions	9
4.1	General description	9
4.2	Numerical example and Monte Carlo exercises	11
4.3	Relationship to other identifying restrictions	13
5.	Uncertainty and Business Cycles Revisited	14
5.1	Identifying restrictions in LMN	14
5.1.1	External variable constraints	14
5.1.2	Event constraints	15
5.1.3	Numerical implementation	15
5.1.4	Results under LMN restrictions	16
5.2	Shock-percentile restrictions	17
5.2.1	Impulse responses	17
5.2.2	How important are uncertainty shocks?	19
5.2.3	Understanding the results	20
5.3	Additional restrictions	21
6.	Estimating the Effects of US Monetary Policy	23
6.1	Reduced-form VAR	24
6.2	Identifying restrictions	24
6.3	Results	25
7.	Conclusion	26
	Appendix A: Analytical Results and Illustration in Bivariate Model	27
	Appendix B: Monte Carlo Exercises	30
	Appendix C: Additional Details about Empirical Exercises	32
	References	35



## 1. Introduction

A vast literature uses structural vector autoregressions (SVARs) to identify the effects of structural shocks. The key challenge in doing this is formulating identifying restrictions that are both credible and sufficiently informative to draw economically useful conclusions. In an important contribution, Ludvigson, Ma and Ng (2017, 2021) propose set identifying SVARs using identifying restrictions that include bounds on shocks in selected episodes – what I refer to as ‘shock-magnitude restrictions’. An example is the restriction that the ‘financial uncertainty shock’ in October 1987 – corresponding to the Black Monday stock market crash – was larger than a specified numeric bound (e.g. four standard deviations). In the context of identifying uncertainty shocks, Ludvigson *et al* (2021) (henceforth, LMN) demonstrate that these restrictions can contribute useful identifying information. A number of papers have adopted shock-magnitude restrictions when identifying uncertainty, and other, shocks.<sup>1</sup>

This paper makes three main contributions. First, I highlight an important feature of the procedure used to specify shock-magnitude restrictions: bounds on shock realisations are typically elicited by inspecting the distribution of shocks induced by simulating values of the SVAR’s structural parameters. I argue that the algorithm used to randomly draw parameters is arbitrary in this setting, which casts doubt on the credibility of these restrictions. Second, I propose a novel type of identifying restriction that is similar in spirit to shock-magnitude restrictions, but that avoids specifying numeric bounds on shocks, which are difficult to credibly elicit. The new ‘shock-percentile restrictions’ impose that the realisation of the shock in a selected episode is more extreme than a specified percentile of the shock’s historical distribution; an example is the restriction that the financial uncertainty shock in the Black Monday episode exceeds the 75th percentile of the historical distribution of financial uncertainty shocks. Third, using these restrictions I revisit the relationship between uncertainty and real activity, and estimate the macroeconomic effects of US monetary policy. Both applications demonstrate that shock-percentile restrictions can provide useful identifying information.

Imposing shock-magnitude restrictions, as in LMN, requires specifying a numeric bound on a shock in a selected period. Eliciting such a bound on the basis of *a priori* or narrative information is difficult; while narrative information may suggest that there was a large shock in some period, it is not necessarily obvious how to translate this into a numeric bound. Instead, LMN propose a simulation-based procedure to elicit a bound. This procedure involves randomly drawing values of the SVAR’s parameters using an algorithm that I describe below. The distribution of parameter draws induces a distribution for the structural shocks given the observed data. LMN then take a specified percentile of the shock distribution in the selected period and use this to impose a shock-magnitude restriction. For example, in the application considered by LMN, this procedure implies that the 75th percentile of the distribution of the financial uncertainty shock in the Black Monday episode is around 4.2 standard deviations. LMN use this result to motivate a shock-magnitude restriction in which the financial uncertainty shock in that episode is constrained to exceed 4.2 standard deviations.

---

1 Several papers employ shock-magnitude restrictions to help identify uncertainty shocks, including Caggiano, Castelnuovo and Kima (2020), Caggiano *et al* (2020), Himounet (2022), Theophilopoulou (2022), Caggiano and Castelnuovo (2023), Carriero, Marcellino and Tornese (2023), Baker, Bloom and Terry (2024) and Kang and Park (2024). Shock-magnitude restrictions have also been used in other settings. For example, they are applied by Ludvigson *et al* (2017) in models of the global oil market and US monetary policy, and by Ferreira (2022) to disentangle forward guidance and conventional monetary policy shocks.

Shock-magnitude restrictions can be interpreted as representing the belief that a shock was ‘large’, where the definition of ‘large’ is elicited based on the simulation-based procedure described above. In this procedure, the algorithm used to randomly draw parameter values is from Rubio-Ramírez, Waggoner and Zha (2010) (henceforth RWZ). The algorithm draws orthonormal (or ‘rotation’) matrices from a uniform distribution and is employed widely when estimating Bayesian SVARs (e.g. Arias, Rubio-Ramírez and Waggoner 2018; Baumeister and Hamilton 2018). In the Bayesian inferential setting, this uniform distribution has a well-defined interpretation as a prior or posterior distribution.<sup>2</sup> However, in the frequentist setting considered in LMN, the distribution induced by this algorithm has no such interpretation; the algorithm is merely a convenient tool for randomly generating draws of the parameters (Fry and Pagan 2011). Consequently, the distribution of the structural shocks induced by this algorithm is arbitrary, which casts doubt on the credibility of shock-magnitude restrictions.<sup>3</sup>

As an alternative, I propose identifying shocks using novel identifying restrictions that are similar in spirit to shock-magnitude restrictions, but that avoid the difficult problem of specifying numeric bounds on shocks. More specifically, I consider imposing that the shock in a selected episode exceeds a specified percentile of the historical distribution of that shock. These shock-percentile restrictions can be used to impose the belief that a shock was large relative to other realisations of the same shock occurring within the sample. Because these restrictions relate the realisation of a shock to a percentile of its historical distribution, the restrictions have a natural interpretation as reflecting beliefs about how ‘rare’ such a shock is, expressed in probabilistic terms. In contrast, shock-magnitude restrictions have no such probabilistic interpretation in the absence of a parametric distributional assumption about the structural shock (e.g. that it is Gaussian).

Shock-percentile restrictions are identifying restrictions that directly involve the realisations of structural shocks in selected episodes. They therefore fall into the broad class of ‘narrative restrictions’ proposed in Antolín-Díaz and Rubio-Ramírez (2018) and analysed in Giacomini, Kitagawa and Read (2023). Shock-percentile restrictions differ to the specific narrative restrictions proposed in Antolín-Díaz and Rubio-Ramírez (2018), which constrain the signs of structural shocks in selected episodes and/or the relative contributions of shocks to forecast errors (i.e. historical decompositions). While restrictions on the relative magnitudes of structural shocks in different periods have been considered elsewhere, these have typically involved restricting a shock to be larger than the realisation of the shock in *all* other periods (e.g. Ben Zeev 2018; Giacomini, Kitagawa and Read 2021; Read 2022; Abbate, Eickmeier and Prieto 2023). I demonstrate that much weaker restrictions on the relative size of structural shocks can, in some cases, contribute substantial identifying power. Using a simple model and Monte Carlo exercises, I illustrate the potential for shock-percentile restrictions to sharpen identification. I then apply shock-percentile restrictions in two empirical applications.

First, I revisit the relationship between uncertainty and real activity by replacing the shock-magnitude restrictions in LMN with shock-percentile restrictions. Doing this yields qualitatively similar conclusions about the effects of *financial* uncertainty shocks on output. In particular, it remains the case that positive financial uncertainty shocks unambiguously lead to a decline in output, and financial uncertainty shocks explain a nontrivial share of output fluctuations beyond short horizons. More

---

2 Arias, Rubio-Ramírez and Waggoner (2025) show that the uniform prior for the orthonormal matrix induces a joint uniform prior over the vector of impulse responses.

3 Section 3.3 argues that a Bayesian analogue of this bound-elicitation procedure would also be problematic.

generally, however, the shock-percentile restrictions tend to yield ambiguous conclusions about the relationship between uncertainty and real activity. Replacing shock-magnitude restrictions with shock-percentile restrictions in this exercise is therefore an example of the trade-off between ‘credibility’ and ‘certitude’ discussed in Manski (2003, 2011).<sup>4</sup> Imposing an additional shock-percentile restriction related to the 1998 Russian Financial Crisis better disentangles shocks to macroeconomic and financial uncertainty and yields results that are qualitatively consistent with the results from LMN across several dimensions.

Second, I exploit a shock-percentile restriction to sharpen identification of the macroeconomic effects of US monetary policy, building on a model from Antolín-Díaz and Rubio-Ramírez (2018). I start with their narrative restrictions, which constrain the sign of the monetary policy shock and its contribution to changes in the federal funds rate in selected episodes. I then impose a single shock-percentile restriction reflecting the belief that the monetary policy shock in October 1979 – corresponding to the Volcker disinflation – was large relative to the historical distribution of monetary policy shocks. Despite the original narrative restrictions already providing substantial identifying power, this single shock-percentile restriction materially sharpens identification of the output responses to a monetary policy shock; relative to the original set of restrictions, there is much stronger evidence that output declines following a positive shock, especially at short horizons.

**Roadmap.** Section 2 describes the SVAR, explains the concept of an identified set and defines shock-magnitude restrictions. Section 3 discusses the procedure used to elicit bounds for shock-magnitude restrictions. Section 4 introduces shock-percentile restrictions, illustrates their potential to contribute identifying power in a simple model and relates them to existing identifying restrictions. Section 5 revisits the relationship between uncertainty and real activity, comparing and contrasting results obtained using shock-magnitude and shock-percentile restrictions. Section 6 exploits a shock-percentile restriction to sharpen identification of the effects of US monetary policy. Section 7 concludes. The appendices contain additional details related to the numerical and empirical exercises.

**Notation.**  $e_{i,n}$  is column  $i$  of the  $n \times n$  identity matrix,  $\mathbf{I}_n$ .  $\mathbf{0}_{n \times m}$  is an  $n \times m$  matrix of zeros. For  $n \times m$  matrix  $\mathbf{X}$ ,  $\text{vec}(\mathbf{X})$  is the vectorisation of  $\mathbf{X}$ , which stacks the elements of  $\mathbf{X}$  into an  $nm \times 1$  vector. If  $\mathbf{X}$  is  $n \times n$ ,  $\text{vech}(\mathbf{X})$  is the half-vectorisation of  $\mathbf{X}$ , which stacks the elements lying on or below the diagonal into an  $n(n+1)/2 \times 1$  vector.  $\mathbb{1}(\cdot)$  is the indicator function.

## 2. Framework

This section describes the SVAR, explains the concept of an identified set and defines shock-magnitude restrictions.

### 2.1 SVAR

Assume  $\mathbf{y}_t = (y_{1t}, \dots, y_{nt})'$  has a reduced-form VAR( $p$ ) representation:

$$\mathbf{y}_t = \sum_{l=1}^p \mathbf{B}_l \mathbf{y}_{t-l} + \mathbf{u}_t \quad (1)$$

---

<sup>4</sup> Shock-percentile restrictions are not necessarily always weaker than comparable shock-magnitude restrictions. See Section 5.2.3 for further discussion.

where  $\mathbf{u}_t$  are serially uncorrelated reduced-form innovations with  $\mathbb{E}(\mathbf{u}_t) = \mathbf{0}_{n \times 1}$  and  $\mathbb{E}(\mathbf{u}_t \mathbf{u}_t') = \mathbf{\Sigma}$ .<sup>5</sup> Let  $\mathbf{\Sigma}_{tr}$  be the lower-triangular Cholesky factor of  $\mathbf{\Sigma}$  with non-negative diagonal elements, so  $\mathbf{\Sigma}_{tr} \mathbf{\Sigma}_{tr}' = \mathbf{\Sigma}$ . Collect the reduced-form parameters into  $\boldsymbol{\phi} = (\text{vec}(\mathbf{B}_1)', \dots, \text{vec}(\mathbf{B}_p)', \text{vech}(\mathbf{\Sigma}_{tr})')'$ .

Assume  $\mathbf{u}_t$  is related to the structural shocks  $\boldsymbol{\varepsilon}_t = (\varepsilon_{1t}, \dots, \varepsilon_{nt})'$  by

$$\mathbf{u}_t = \mathbf{H} \boldsymbol{\varepsilon}_t \quad (2)$$

where  $\mathbb{E}(\boldsymbol{\varepsilon}_t) = \mathbf{0}_{n \times 1}$ ,  $\mathbb{E}(\boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_t') = \mathbf{I}_n$  and  $\mathbf{H}$  is invertible.  $\mathbf{H}$  contains the impact impulse responses of the variables in  $\mathbf{y}_t$  to the structural shocks. The diagonal elements of  $\mathbf{H}$  are normalised to be non-negative (a 'sign normalisation'), so a positive realisation of structural shock  $i$  does not decrease variable  $i$  on impact. Knowledge of  $\mathbf{H}$  and the reduced-form parameters  $\boldsymbol{\phi}$  allows impulse responses and forecast error variance decompositions (FEVDs) to be computed at any horizon.<sup>6</sup>

## 2.2 Identified sets

In set-identified SVARs, it is common to work with the model's 'orthogonal reduced-form' parameterisation (e.g. Arias *et al* 2018). Let  $\mathbf{H} = \mathbf{\Sigma}_{tr} \mathbf{Q}$ , where  $\mathbf{Q} \in \mathcal{O}(n)$  is a  $n \times n$  orthonormal matrix and  $\mathcal{O}(n)$  is the space of all such matrices. In the absence of further identifying restrictions,  $\mathbf{Q}$  is set identified and, consequently, so is  $\mathbf{H}$  (e.g. Uhlig 2005; RWZ).

Given a value of  $\boldsymbol{\phi}$ , define the 'unconstrained identified set' for  $\mathbf{Q}$  as:

$$\mathcal{Q}(\boldsymbol{\phi}) = \{\mathbf{Q} \in \mathcal{O}(n) : \text{diag}(\mathbf{\Sigma}_{tr} \mathbf{Q}) \geq \mathbf{0}_{n \times 1}\} \quad (3)$$

Values of  $\mathbf{Q}$  in  $\mathcal{Q}(\boldsymbol{\phi})$  are observationally equivalent, in the sense that they correspond to the same second moments of the data, which are summarised by  $\boldsymbol{\phi}$ .  $\mathcal{Q}(\boldsymbol{\phi})$  induces an unconstrained identified set for any object of interest that is a function of  $\mathbf{Q}$ . For example, the unconstrained identified set for  $\mathbf{H}$  is  $\{\mathbf{H} = \mathbf{\Sigma}_{tr} \mathbf{Q} : \mathbf{Q} \in \mathcal{Q}(\boldsymbol{\phi})\}$ . Unconstrained identified sets completely characterise the information contained in the data about the objects of interest when imposing only the sign normalisation.

Traditionally, sign or zero restrictions would be imposed on functions of  $\mathbf{H}$ , shrinking the identified set, with point identification attained if there are sufficiently many zero restrictions. However, as argued in LMN in the context of identifying uncertainty shocks, such restrictions have little theoretical motivation given the ambiguous predictions of different theories of uncertainty.<sup>7</sup> LMN therefore turn to alternative identifying restrictions. These include 'event constraints', which are inequality restrictions on structural shocks in selected episodes. I describe these identifying restrictions in detail below. For now, it suffices to note that these identifying restrictions can be written as inequality restrictions on  $\mathbf{Q}$ . These restrictions can truncate the identified set for  $\mathbf{Q}$ , sharpening identification of the parameters of interest, but in general leaving them set identified.

Given a generic set of  $s$  identifying restrictions,  $S(\boldsymbol{\phi}, \mathbf{Q}) \geq \mathbf{0}_{s \times 1}$ , the identified set for  $\mathbf{Q}$  is:

$$\mathcal{Q}(\boldsymbol{\phi} | S) = \{\mathbf{Q} \in \mathcal{O}(n) : \text{diag}(\mathbf{\Sigma}_{tr} \mathbf{Q}) \geq \mathbf{0}_{n \times 1}, S(\boldsymbol{\phi}, \mathbf{Q}) \geq \mathbf{0}_{s \times 1}\} \quad (4)$$

<sup>5</sup> I abstract from the inclusion of a constant for simplicity.

<sup>6</sup> For definitions of impulse responses and FEVDs, see Kilian and Lütkepohl (2017).

<sup>7</sup> Similar arguments have been made in other settings. For example, Ramey (2016) argues that the recursive orderings commonly used when identifying the effects of monetary policy have little theoretical motivation.

In the discussion below, it will become apparent that event constraints depend on additional objects not captured in the definition of  $\phi$ , but I suppress this dependence for simplicity.<sup>8</sup> The identified set for  $\mathbf{Q}$  induces identified sets for other parameters, such as impulse responses and FEVDs.

### 2.3 Shock-magnitude restrictions

The event constraints proposed in LMN restrict the values of the structural shocks in selected historical episodes. Since  $\boldsymbol{\varepsilon}_t = \boldsymbol{\varepsilon}_t(\phi, \mathbf{Q}, \mathbf{u}_t) = \mathbf{Q}'\boldsymbol{\Sigma}_{tr}^{-1}\mathbf{u}_t$ , the event constraints represent restrictions on  $\mathbf{Q}$  that depend on the reduced-form innovations in the constrained periods in addition to the reduced-form parameters  $\phi$ .

LMN consider different types of event constraints, which I describe in Section 5. The focus of the following discussion is on the event constraints that involve imposing a bound on a shock in a selected period, which I refer to as shock-magnitude restrictions. A shock-magnitude restriction on shock  $j$  in period  $t$  takes the form

$$\varepsilon_{jt} = \mathbf{e}'_{j,n}\mathbf{Q}'\boldsymbol{\Sigma}_{tr}^{-1}\mathbf{u}_t = (\boldsymbol{\Sigma}_{tr}^{-1}\mathbf{u}_t)'\mathbf{q}_j \geq k \quad (5)$$

where  $k \in \mathbb{R}$  is a numeric lower bound (specified by the researcher) and  $\mathbf{q}_j$  is column  $j$  of  $\mathbf{Q}$ .<sup>9</sup> An upper bound would be imposed similarly.

An important input into the shock-magnitude restriction is the bound,  $k$ . More-extreme values of  $k$  make the restriction more likely to bind and thus truncate the identified set for  $\mathbf{Q}$ , potentially delivering sharper identification of the SVAR's structural parameters.<sup>10</sup> The next section discusses the procedure proposed in LMN to elicit bounds.

## 3. Challenges with Eliciting Bounds on Shocks

Imposing a shock-magnitude restriction requires specifying the numeric bound,  $k$ , in (5) above. While narrative information may suggest that a shock was large in a specific episode, it is difficult to use this information to formulate a numeric bound on the shock. If we knew that the structural shocks were Gaussian, then restricting a shock to exceed, say, three standard deviations in size could be interpreted as imposing the belief that the period contained a shock with magnitude that is expected to occur with low probability. However, at this point, the SVAR framework considered here makes no parametric distributional assumption about the shocks. It is therefore not obvious whether a shock exceeding three standard deviations is a low-probability event before observing the data and identifying the structural shocks.

Instead, LMN propose a procedure for eliciting bounds on shock magnitudes that uses the observed data and its estimated covariance structure (i.e. the VAR's reduced-form parameters) as inputs. An important additional ingredient in this procedure is an algorithm that is used to simulate values of  $\mathbf{Q}$ . Section 3.1 describes this algorithm. Section 3.2 then argues that the arbitrary nature of the algorithm

<sup>8</sup> When the identifying restrictions include event constraints (or narrative restrictions), the restrictions depend on the realisation of the data via the reduced-form innovations,  $\mathbf{u}_t$ . This means the standard definition of an identified set as the set of observationally equivalent parameters does not apply. In settings like this, Giacomini *et al* (2023) introduce a refinement of the identified set – a 'conditional identified set' – that describes the data-dependent mapping from reduced-form parameters to the structural object of interest. For ease of exposition, I do not distinguish between conditional identified sets and identified sets.

<sup>9</sup> Given the normalisation that  $\mathbb{E}(\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t') = \mathbf{I}_n$ ,  $k$  is measured in standard deviations.

<sup>10</sup> Ludvigson, Ma and Ng (2017) discuss how shock-magnitude restrictions can provide identifying information.

in this setting casts doubt on the credibility of the resulting shock-magnitude restrictions. Section 3.3 argues that the procedure would also be problematic if viewed from a Bayesian perspective.

### 3.1 The RWZ algorithm

A key ingredient in the simulation-based procedure used to elicit the shock-magnitude bounds is an algorithm used to obtain random draws of  $\mathbf{Q}$ . The algorithm, based on Rubio-Ramírez *et al* (2010), involves the following steps:

1. Draw an  $n \times n$  matrix of independent standard normal random variables  $\mathbf{Z}$ .
2. Compute the QR decomposition  $\mathbf{Z} = \mathbf{QR}$ , where  $\mathbf{R}$  is upper-triangular with non-negative diagonal elements.
3. Normalise the columns of  $\mathbf{Q}$  to satisfy  $\text{diag}(\mathbf{H}) = \text{diag}(\mathbf{\Sigma}_{rr}\mathbf{Q}) \geq \mathbf{0}_{n \times 1}$ .

The resulting draws of  $\mathbf{Q}$  are uniformly distributed over the unconstrained identified set  $\mathcal{Q}(\phi)$ .

### 3.2 The bound elicitation procedure

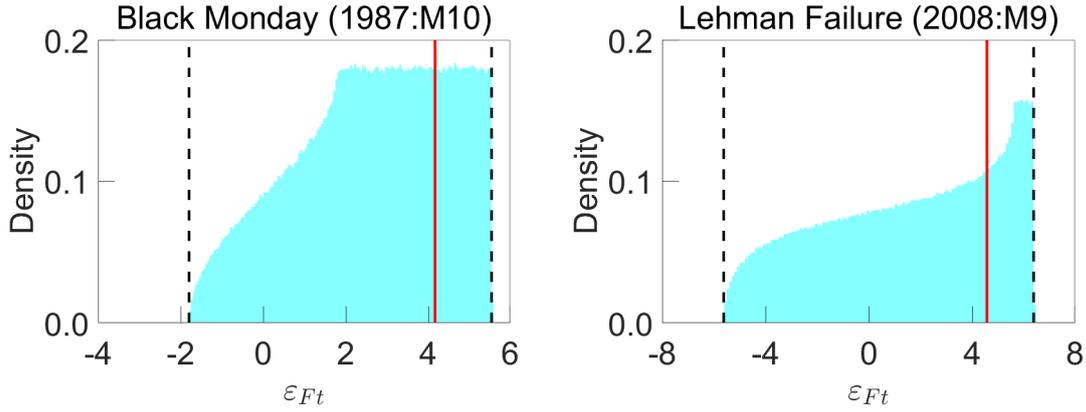
In illustrating the procedure proposed by LMN to elicit bounds on shock magnitudes, it is useful to introduce their empirical setting. They use a three-variable SVAR to examine the relationship between real activity and measures of macroeconomic and financial uncertainty in the United States.

Let  $\mathbf{y}_t = (M_t, Y_t, F_t)'$  with corresponding structural shocks  $\boldsymbol{\varepsilon}_t = (\varepsilon_{M_t}, \varepsilon_{Y_t}, \varepsilon_{F_t})'$ .  $M_t$  is the measure of macroeconomic uncertainty constructed in Jurado, Ludvigson and Ng (2015). This measure averages the square root of conditional forecast error variances of one-step-ahead forecasts across a large panel of macroeconomic data, where the conditional variances are estimated using stochastic volatility models. Underlying this measure is the idea that macroeconomic uncertainty is higher when macroeconomic variables are more difficult to predict.  $Y_t$  is the log of industrial production, which is a measure of real activity.  $F_t$  is a measure of financial uncertainty constructed in the same way as  $M_t$ , but based on a large set of financial indicators. The VAR is monthly and contains six lags and a constant. The reduced-form parameters  $\phi$  are estimated via ordinary least squares (OLS) over a sample running from July 1960 to April 2015.<sup>11</sup>

LMN use the RWZ algorithm to obtain 1.5 million draws of  $\mathbf{Q}$  from  $\mathcal{Q}(\phi)$ . Each draw implies a path for the structural shocks,  $\{\boldsymbol{\varepsilon}_t(\phi, \mathbf{Q}, \mathbf{u}_t)\}_{t=1}^T$ . Computing the sequence of shocks at each draw of  $\mathbf{Q}$  builds up a shock distribution in each period. To motivate their shock-magnitude restrictions, LMN examine these shock distributions. Figure 1 plots the distributions of the 'financial uncertainty shock'  $\varepsilon_{F_t}(\phi, \mathbf{Q}, \mathbf{u}_t)$  in September 2008 (the Lehman failure) and October 1987 (the Black Monday episode). In both episodes, the distribution of the shock is right skewed with positive values more likely than negative values.

---

<sup>11</sup> Angelini *et al* (2019) consider the same model, but identify shocks by exploiting breaks in the volatility of macroeconomic variables.

**Figure 1: Distribution of Financial Uncertainty Shocks Implied by RWZ Algorithm**

Notes: Distributions of financial uncertainty shocks implied by RWZ algorithm when imposing sign normalisation on impact impulse responses. Solid lines are 75th percentiles; dashed lines are bounds of identified sets.

LMN interpret these results as suggesting that ‘the covariance structure of the data alone provides overwhelming evidence of a large positive financial uncertainty shock in the months of the 1987 crash and the 2008 Lehman collapse’. A similar exercise for the macroeconomic uncertainty shock points to large positive macroeconomic uncertainty shocks in September 2008 (the Lehman episode) and December 1970 (leading up to the collapse of the Bretton Woods system of fixed exchange rates). LMN subsequently formulate shock-magnitude restrictions in these periods, where the bounds ( $k$ ) are elicited based on these shock distributions. For instance, one exercise in LMN sets lower bounds equal to the 75th percentiles of the shock distributions. To give a specific example, the financial uncertainty shock in the Black Monday episode is restricted to exceed the 75th percentile of the shock distribution in that period, which is around 4.2 standard deviations; that is, the shock-magnitude restriction is  $\varepsilon_{Ft}(\phi, \mathbf{Q}, \mathbf{u}_t) \geq 4.2$  for  $t = 1987:M10$ .

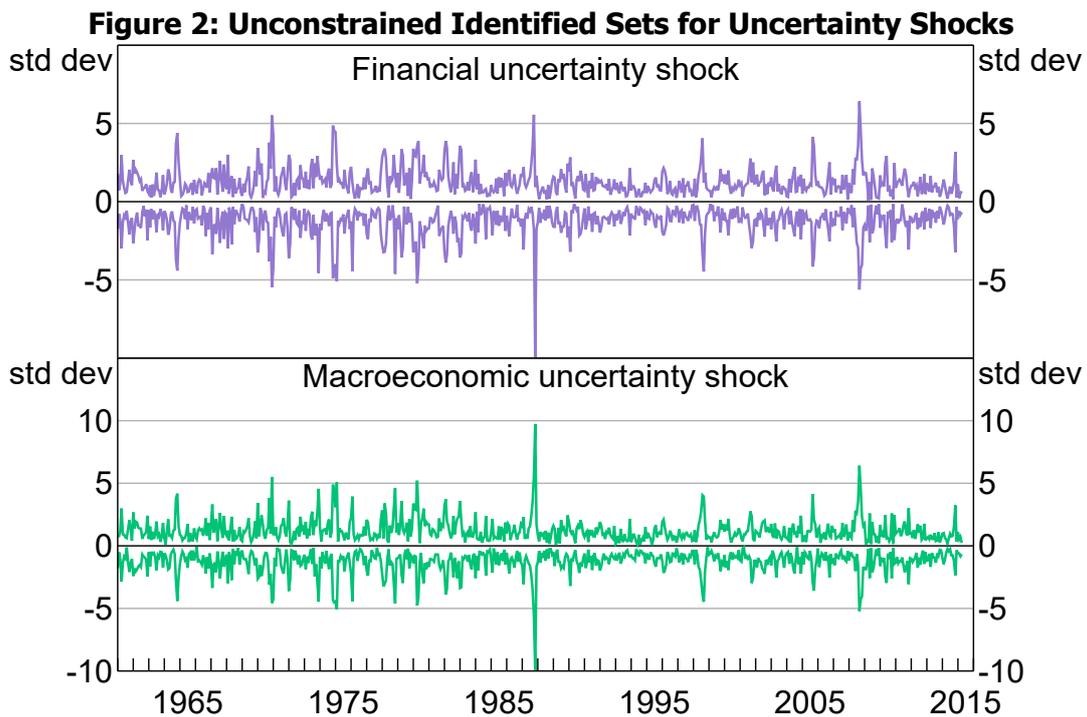
This procedure has been used by several papers that employ shock-magnitude restrictions (see the references in footnote 1). However, given the sign normalisation, the structural parameters are set identified and any  $\mathbf{Q}$  in  $\mathcal{Q}(\phi)$  is observationally equivalent. The distribution for  $\mathbf{Q}$  induced by the RWZ (or any other) algorithm is therefore arbitrary in this setting (e.g. Fry and Pagan 2011; Baumeister and Hamilton 2015, 2024). In turn, the shock distributions induced by the distribution for  $\mathbf{Q}$  are also arbitrary. Ultimately, this casts doubt on the credibility of shock-magnitude restrictions when the bounds are elicited via this procedure.

To support the claim that the induced shock distribution is arbitrary, I examine unconstrained identified sets for the uncertainty shocks in each month:<sup>12</sup>

$$\{\varepsilon_{jt}(\phi, \mathbf{Q}, \mathbf{u}_t) : \mathbf{Q} \in \mathcal{Q}(\phi)\} = \left[ \inf_{\mathbf{Q} \in \mathcal{Q}(\phi)} \varepsilon_{jt}(\phi, \mathbf{Q}, \mathbf{u}_t), \sup_{\mathbf{Q} \in \mathcal{Q}(\phi)} \varepsilon_{jt}(\phi, \mathbf{Q}, \mathbf{u}_t) \right] \quad (6)$$

<sup>12</sup> These identified sets can be characterised by their end points.  $\varepsilon_{jt}(\phi, \mathbf{Q}, \mathbf{u}_t) = (\boldsymbol{\Sigma}_{tr}^{-1} \mathbf{u}_t)' \mathbf{q}_j$  depends on  $\mathbf{Q}$  only via  $\mathbf{q}_j$ . The unconstrained identified set for  $\mathbf{q}_j$  is given by the intersection of the unit sphere in  $\mathbb{R}^3$  and the half-space  $\mathbf{e}'_{j,3} \boldsymbol{\Sigma}_{tr} \mathbf{q}_j \geq 0$ , which is path-connected. Since  $\varepsilon_{jt}(\phi, \mathbf{Q}, \mathbf{u}_t)$  is a continuous function of  $\mathbf{q}_j$  on a path-connected domain, the set of function values on this domain is an interval. This argument follows the same reasoning as in Giacomini and Kitagawa (2021a, 2021b), who provide conditions for convexity of impulse-response identified sets. I approximate these sets by taking the minimum and maximum of  $\varepsilon_{jt}(\phi, \mathbf{Q}, \mathbf{u}_t)$  over the 1.5 million draws of  $\mathbf{Q}$  from  $\mathcal{Q}(\phi)$ .

These unconstrained identified sets characterise the information about the structural shocks contained in the data, given the sign normalisation. Figure 2 plots the identified sets for the uncertainty shocks in each period. It is apparent that the identified sets for the shocks always span zero, so – given the sign normalisation only – the data and its covariance structure never unambiguously imply whether a positive or negative shock has occurred, much less whether the shock was relatively large or small. At best, we may be able to rule out certain features of the shocks; for example, the identified set for the financial uncertainty shock in the Black Monday episode has a lower bound of around  $-2$  (see Figure 1), so the data are inconsistent with a very large negative financial uncertainty shock in this period. Otherwise, the features of the shock distributions used to elicit the shock-magnitude restrictions depend on the algorithm used to draw  $\mathbf{Q}$ . If we used a different algorithm, we could obtain a very different distribution.



The unconstrained identified sets for the two uncertainty shocks include zero in every period and the data are largely uninformative about the magnitudes of the shocks. A natural question is whether this result is specific to the current application – for example, does it depend on features of the reduced-form parameters? Appendix A.1 shows that, in a general  $n$ -variable SVAR, unconstrained identified sets for structural shocks include zero at all possible values of the reduced-form parameters and for any realisation of the data. So the indeterminacy of the shock signs is not peculiar to the current application. Appendix A.2 uses a simple example to illustrate how different algorithms for randomly generating parameters could in principle generate very different shock distributions.

### 3.3 Can the procedure be motivated from a Bayesian perspective?

In the Bayesian setting, the RWZ algorithm is often used to draw from the posterior distribution under a uniform prior for  $\mathbf{Q}$ .<sup>13</sup> A Bayesian analogue of the bound-elicitation procedure would be to: 1) specify a prior for  $\phi$  and a uniform prior over  $\mathcal{Q}(\phi)$ ; 2) use the posterior for the structural shocks to elicit bounds (e.g. by taking a specified percentile of the posterior distributions of the shocks in selected periods); and 3) use these bounds to impose shock-magnitude restrictions in another SVAR.

There are two main problems with this procedure. First, the obtained bounds may be sensitive to the choice of prior assumed in the first step; since  $\mathbf{Q}$  is set identified, the conditional prior for  $\mathbf{Q}$  given  $\phi$  is not updated by the likelihood (e.g. Poirier 1998; Moon and Schorfheide 2012). To the extent that the uniform prior does not accurately reflect the researcher's beliefs, this sensitivity is problematic. Second, even if the uniform prior for  $\mathbf{Q}$  is a good representation of the researcher's subjective beliefs, it is not obvious how one could justify dogmatic identifying restrictions on the basis of the induced posterior over the structural shocks; the posterior merely tells us that some shock values are more or less likely than others, whereas a shock-magnitude restriction dogmatically rules out some shock values. In summary, even with a well-defined subjective (uniform) prior for  $\mathbf{Q}$ , a Bayesian analogue of the bound-elicitation procedure would generate shock-magnitude restrictions that lack credibility.

## 4. Shock-percentile Restrictions

This section proposes alternative identifying restrictions that are similar in spirit to shock-magnitude restrictions, but that avoid the difficult problem of eliciting credible numeric bounds on shocks.

### 4.1 General description

Shock-magnitude restrictions can be viewed as imposing the belief that shocks in selected episodes were 'large'. Rather than eliciting the definition of 'large' based on the distribution of shocks induced by the RWZ algorithm, a natural alternative is to define 'large' as being relative to a percentile of the *historical* distribution of shocks. I refer to such restrictions as 'shock-percentile restrictions'.

More specifically, let

$$F_j(x; \phi, \mathbf{Q}, \mathbf{U}) = (1/T) \sum_{t=1}^T \mathbb{1}(\varepsilon_{jt}(\phi, \mathbf{Q}, \mathbf{u}_t) \leq x) \quad (7)$$

be the cumulative distribution function (CDF) of the shocks  $\{\varepsilon_{jt}(\phi, \mathbf{Q}, \mathbf{u}_t)\}_{t=1}^T$ , evaluated at  $x \in \mathbb{R}$  and where  $\mathbf{U} = (\mathbf{u}'_1, \dots, \mathbf{u}'_T)'$ . For  $\alpha \in [0, 1]$ , the quantile function  $G_j(\alpha; \phi, \mathbf{Q}, \mathbf{U})$  satisfies

$$G_j(\alpha; \phi, \mathbf{Q}, \mathbf{U}) = \inf_x \{x \in \mathbb{R} : F_j(x; \phi, \mathbf{Q}, \mathbf{U}) \geq \alpha\} \quad (8)$$

A shock-percentile restriction on shock  $j$  in period  $\tau$  is then

$$\varepsilon_{j\tau}(\phi, \mathbf{Q}, \mathbf{u}_\tau) \geq G_j(\alpha; \phi, \mathbf{Q}, \mathbf{U}) \quad (9)$$

The restriction that shock  $j$  in period  $\tau$  is smaller than a specified percentile is defined similarly. These restrictions have a natural interpretation as reflecting beliefs about how extreme the shock in

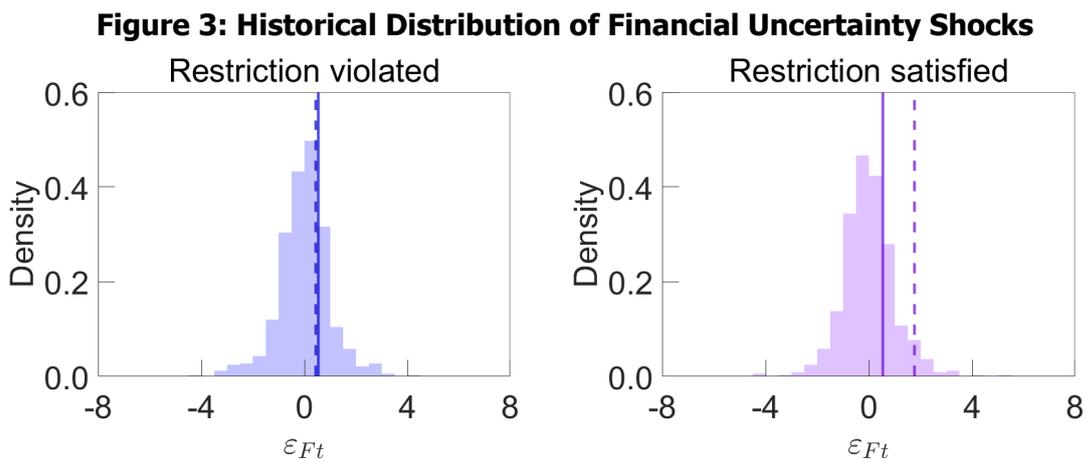
---

<sup>13</sup> While there has been debate about the desirability of using such a prior, it nevertheless remains a well-defined prior distribution. For different sides of this debate, see Baumeister and Hamilton (2018, 2019, 2024), Watson (2020), Arias *et al* (2025) and Inoue and Kilian (2025, 2026).

a selected episode was, expressed as a relative frequency. To give an example in the context of LMN's model of uncertainty, setting  $\tau = 1987:M10$ ,  $j = F$  and  $\alpha = 0.75$  requires the financial uncertainty shock in the Black Monday episode to exceed the 75th percentile of the historical distribution of financial uncertainty shocks; in other words, (positive) shocks more extreme than the shock in this episode occurred in no more than 25 per cent of periods within the sample. In contrast, shock-magnitude restrictions have no such interpretation in the absence of a parametric distributional assumption (e.g. that the shocks are Gaussian).

Shock-percentile restrictions can be used to impose the belief that shocks in selected episodes were large relative to their historical distributions. Importantly, imposing this belief does not require specifying a numeric bound, which as argued above is difficult to credibly elicit. Instead, given a choice of  $\alpha$ , the bound is a function of the historical distribution of realised shocks. Because the shocks are in general set identified, we do not know the exact realisations of the shocks nor the percentiles of their historical distribution; different values of  $\mathbf{Q}$  will imply different shock sequences. Some shock sequences may be inconsistent with the shock-percentile restrictions. By ruling out parameter values implying shock sequences where the shock-percentile restrictions are violated, these restrictions may truncate the identified set for  $\mathbf{Q}$ , sharpening identification.

To illustrate this idea, consider again the SVAR from LMN. Fixing  $\phi$  at the OLS estimates, different values of  $\mathbf{Q}$  will imply different sequences for the financial uncertainty shock  $\varepsilon_{Ft}$ . The left panel of Figure 3 plots the distribution of  $\{\varepsilon_{Ft}(\phi, \mathbf{Q}, \mathbf{u}_t)\}_{t=1}^T$  at a random value of  $\mathbf{Q}$  in the unconstrained identified set  $\mathcal{Q}(\phi)$ . In this case, the value of  $\varepsilon_{Ft}(\phi, \mathbf{Q}, \mathbf{u}_t)$  in the Black Monday episode (the dashed line) lies below the 75th percentile of the historical distribution of  $\varepsilon_{Ft}(\phi, \mathbf{Q}, \mathbf{u}_t)$  (the solid line). This value of  $\mathbf{Q}$  would therefore be ruled out by the shock-percentile restriction  $\varepsilon_{F\tau}(\phi, \mathbf{Q}, \mathbf{u}_\tau) \geq G_F(0.75; \phi, \mathbf{Q}, \mathbf{U})$  (for  $\tau = 1987:M10$ ). In the right panel, which considers a different value of  $\mathbf{Q}$ , the realisation of  $\varepsilon_{F\tau}(\phi, \mathbf{Q}, \mathbf{u}_\tau)$  in the Black Monday episode exceeds the 75th percentile of the historical distribution of  $\varepsilon_{Ft}(\phi, \mathbf{Q}, \mathbf{u}_t)$ , so this value of  $\mathbf{Q}$  would be retained within the identified set given this shock-percentile restriction.



Notes: Each panel represents historical distribution of  $\varepsilon_{Ft}(\phi, \mathbf{Q}, \mathbf{u}_t)$  evaluated at a different (random) value of  $\mathbf{Q}$ . Solid lines represent 75th percentile of distribution; dashed lines represent value of  $\varepsilon_{F\tau}(\phi, \mathbf{Q}, \mathbf{u}_\tau)$  at  $\tau = 1987:M10$ .

Shock-percentile restrictions do not necessarily translate into numeric bounds on shock magnitudes; since the shocks are set identified, in general so is  $G_j(\alpha; \phi, \mathbf{Q}, \mathbf{U})$ .<sup>14</sup> One case where shock-percentile restrictions can be reformulated as shock-magnitude restrictions is when the shocks are Gaussian. In this case, the population  $\alpha$  percentiles of the marginal shock distributions are  $\Phi^{-1}(\alpha)$  where  $\Phi^{-1}(\cdot)$  is the standard normal inverse CDF. Setting  $\phi$  to its true value  $\phi_0$ ,  $G_j(\alpha; \phi_0, \mathbf{Q}, \mathbf{U}) \xrightarrow{P} \Phi^{-1}(\alpha)$  for any  $\mathbf{Q} \in \mathcal{Q}(\phi)$ .<sup>15</sup> In this case, the unconstrained identified set for  $G_j(\alpha; \phi, \mathbf{Q}, \mathbf{U})$  would collapse to a point at  $\Phi^{-1}(\alpha)$ , and a shock-percentile restriction would resemble a numeric lower bound on the shock; imposing the shock-percentile restriction (9) would be (asymptotically) equivalent to the shock-magnitude restriction  $\varepsilon_{j\tau}(\phi, \mathbf{Q}, \mathbf{u}_\tau) \geq \Phi^{-1}(\alpha)$ . Regardless of whether the underlying shocks are Gaussian, in which case the percentiles of the historical shock distribution are point identified, or non-Gaussian, in which case the historical percentiles may be set identified, shock-percentile restrictions can be used to sharpen identification.<sup>16</sup>

## 4.2 Numerical example and Monte Carlo exercises

This section illustrates the potential for shock-percentile restrictions to sharpen identification using the simplest possible SVAR – a bivariate SVAR(0).

Let  $\mathbf{y}_t = (p_t, q_t)'$ , where  $p_t$  is log price and  $q_t$  is log quantity. Assume  $\mathbf{y}_t$  is generated according to:

$$\mathbf{y}_t = \mathbf{H}\boldsymbol{\varepsilon}_t \quad (10)$$

where  $\boldsymbol{\varepsilon}_t \sim N(\mathbf{0}_{2 \times 1}, \mathbf{I}_2)$ .<sup>17</sup> Assume a baseline set of sign restrictions on  $\mathbf{H}$  satisfying

$$\mathbf{H} = \begin{bmatrix} + & + \\ - & + \end{bmatrix} \quad (11)$$

so the first shock can be interpreted as a supply shock that moves prices and quantities in opposite directions, and the second shock can be interpreted as a demand shock that moves prices and quantities in the same direction. The space of  $2 \times 2$  orthonormal matrices can be represented as

$$\mathcal{O}(2) = \left\{ \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \right\} \cup \left\{ \begin{bmatrix} \cos \theta & \sin \theta \\ \sin \theta & -\cos \theta \end{bmatrix} \right\} \quad (12)$$

where  $\theta \in [-\pi, \pi]$  (e.g. Baumeister and Hamilton 2015). The sign restrictions on  $\mathbf{H}$  generate an identified set for  $\theta$ ,  $IS_\theta(\phi) = \{\theta : \mathbf{Q}(\theta) \in \mathcal{Q}(\phi|S)\}$ , which is an interval whose end points depend on  $\phi$ .<sup>18</sup> The scalar  $\theta$  summarises the set-identified component of the model's structural parameters and can be used to illustrate the degree to which the shock-percentile restrictions sharpen identification.

<sup>14</sup> In the example above, the unconstrained identified set for  $G_F(\alpha; \phi, \mathbf{Q}, \mathbf{U})$  is  $[0.43, 0.64]$ .

<sup>15</sup> If  $\boldsymbol{\varepsilon}_t \sim N(\mathbf{0}_{n \times 1}, \mathbf{I}_n)$ , then  $\mathbf{u}_t \sim N(\mathbf{0}_{n \times 1}, \boldsymbol{\Sigma}_0)$ , where  $\boldsymbol{\Sigma}_0 = \mathbf{H}_0 \mathbf{H}_0'$  and  $\mathbf{H}_0$  is the true value of  $\mathbf{H}$ . Since the reduced-form parameters are point identified, we can consider setting  $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_0$ , and  $\boldsymbol{\Sigma}_{rr}^{-1} \mathbf{u}_t \sim N(\mathbf{0}_{n \times 1}, \mathbf{I}_n)$ . Since  $\mathbf{Q}$  is orthonormal, it follows that  $\boldsymbol{\varepsilon}_t(\phi_0, \mathbf{Q}, \mathbf{u}_t) = \mathbf{q}'_j \boldsymbol{\Sigma}_{rr}^{-1} \mathbf{u}_t \sim N(\mathbf{0}_{n \times 1}, \mathbf{I}_n)$  for any  $\mathbf{Q} \in \mathcal{O}(n)$ .

<sup>16</sup> For approaches to identifying SVARs that exploit non-Gaussianity of shocks see, for example, Lanne, Meitz and Saikkonen (2017), Braun (2023) and Andrade, Ferroni and Melosi (2024).

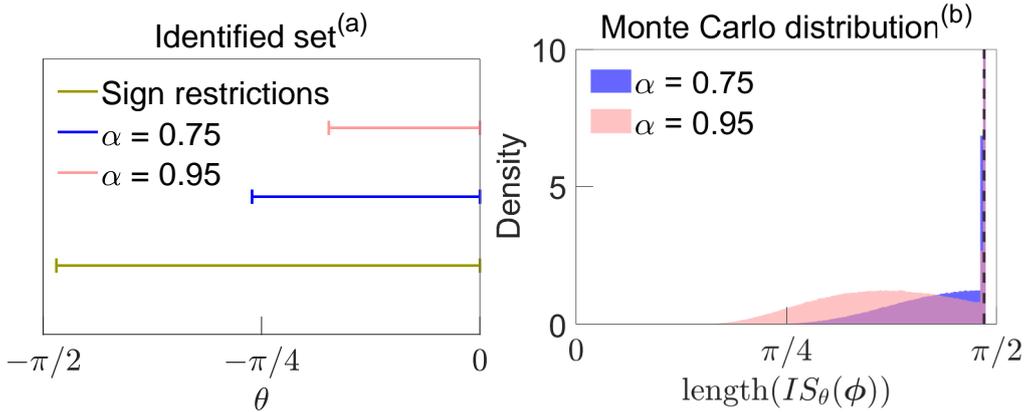
<sup>17</sup> I assume Gaussian shocks here for simplicity, but none of the qualitative results substantively depend on this assumption.

<sup>18</sup> See Baumeister and Hamilton (2015) or Read (2024) for a characterisation of this identified set.

Given a sequence of structural shocks  $\{\boldsymbol{\varepsilon}_t\}_{t=1}^T$  and data  $\{\mathbf{y}_t\}_{t=1}^T$ , impose a single shock-percentile restriction on the supply shock  $\varepsilon_{1\tau}$  in period  $\tau = \arg \max_t \varepsilon_{1t}$ .<sup>19</sup> For the purposes of illustration, the shock-percentile restriction is imposed for different values of  $\alpha \in \{0.75, 0.95\}$ . The premise of this exercise is that we know the period in which a large supply shock has occurred and impose as an identifying restriction that the shock is large relative to the historical distribution of supply shocks. The exercises here assume that  $T = 600$ , which corresponds to 50 years of monthly data. For simplicity, I assume  $\boldsymbol{\phi}$  is known and so abstract from sampling uncertainty.

The left panel of Figure 4 plots the identified set for  $\theta$  under a combination of the sign restrictions on  $\mathbf{H}$  and the shock-percentile restrictions. When imposing the sign restrictions only, the identified set for  $\theta$  is  $[\arctan(\sigma_{22}/\sigma_{21}), 0]$ , where  $\sigma_{22} = \mathbf{e}'_{2,2}\boldsymbol{\Sigma}_{tr}\mathbf{e}_{2,2}$  and  $\sigma_{21} = \mathbf{e}'_{2,2}\boldsymbol{\Sigma}_{tr}\mathbf{e}_{1,2}$ . Additionally imposing the shock-percentile restriction with  $\alpha = 0.75$  narrows the identified set relative to the case where only the sign restrictions are imposed, ruling out lower values of  $\theta$ . Imposing the stronger shock-percentile restriction with  $\alpha = 0.95$  shrinks the identified set further. This illustration demonstrates that the shock-percentile restrictions may sharpen identification conditional on a single realisation from a data-generating process (DGP).

**Figure 4: Illustration of Shock-percentile Restrictions – Bivariate Model**



Notes: (a) Identified set for  $\theta$ ,  $IS_\theta(\boldsymbol{\phi})$ , under different identifying restrictions, conditional on single realisation from DGP.  $\alpha = 0.75$  and  $\alpha = 0.95$  add shock-percentile restrictions to sign restrictions on  $\mathbf{H}$ . (b) Monte Carlo distribution of length of identified set under combination of baseline and shock-percentile restrictions, where each Monte Carlo replication represents different realisation from DGP. Dashed line is length of identified set under baseline sign restrictions. y-axis truncated for readability.

In the exercise above, the shock-percentile restrictions bind and thus contribute identifying information. However, this is not necessarily the case – whether the restriction binds depends on the realisation of the data. To illustrate this, I repeat this exercise at  $10^6$  random realisations of the data and compute the length of  $IS_\theta(\boldsymbol{\phi})$ . The right panel of Figure 4 plots the Monte Carlo distribution of this length.<sup>20</sup> Under the weaker shock-percentile restriction ( $\alpha = 0.75$ ), the identified set is about 7 per cent shorter on average than under the baseline sign restrictions. Even though this shock-percentile restriction tends to sharpen identification on average, in around half of the Monte Carlo replications the restriction contributes no additional identifying information. In contrast, the stronger

<sup>19</sup> The data-generating process underlying this exercise assumes that  $\text{vec}(\mathbf{H}) = (1, -0.3, 0.2, 1.2)'$ , which implies that  $p_t$  and  $q_t$  are weakly negatively correlated.

<sup>20</sup> See Appendix B.1 for details about how this Monte Carlo exercise is implemented.

shock-percentile restriction ( $\alpha = 0.95$ ) reduces the length of the identified set by about one-fifth on average and is binding in around 80 per cent of replications.

In this example, the extent to which the shock-percentile restriction sharpens identification depends on how tight identification is under the baseline sign restrictions; when the baseline identified set is already narrow, the restriction is unlikely to bind and the expected reduction in the set's length is small. This example also imposes only a single shock-percentile restriction; imposing multiple restrictions will tend to sharpen identification further. See Appendix B.2 for illustrations of these cases.

### 4.3 Relationship to other identifying restrictions

Traditional sign restrictions, such as on impulse responses (e.g. Uhlig 2005), directly restrict the SVAR's structural parameters. In contrast, shock-percentile restrictions are restrictions on the structural *shocks*. They can therefore be viewed as falling within the broader class of 'narrative restrictions' proposed in Antolín-Díaz and Rubio-Ramírez (2018) and examined in Giacomini *et al* (2023).<sup>21</sup> The basic premise of narrative restrictions is that, by forcing the model-implied structural shocks to be consistent with narratives about selected historical episodes, they can constrain the structural parameters and hence sharpen identification.

The shock-percentile restrictions proposed here are novel in that they explicitly involve percentiles of historical shock distributions. However, the shock-percentile restrictions possess similarities to some existing applications of narrative restrictions. When identifying investment-specific technology news shocks, Ben Zeev (2018) restricts the maximal three-year average of the shock to occur during the 1997–1999 period, corresponding to the dot-com boom. When identifying US monetary policy shocks, Giacomini *et al* (2021) and Read (2022) consider the restriction that the largest absolute realisation of the shock during their sample period occurred in October 1979, corresponding to the Volcker episode. When identifying financial shocks, Abbate *et al* (2023) impose that the largest contractionary financial shock over their sample period occurred in September quarter 2008, corresponding to the failure of Lehman Brothers. These restrictions have involved restricting a shock in a selected episode to be larger than the realisation of the shock in *all* other periods. They can therefore be viewed as extreme cases of the shock-percentile restrictions proposed here.<sup>22</sup>

Shock-percentile restrictions differ conceptually to narrative restrictions on the historical decomposition, which were proposed by Antolín-Díaz and Rubio-Ramírez (2018). For example, they restrict the contribution of a shock to the change in a variable over some period to be larger than the contributions of other shocks. These kinds of restrictions assign a larger or smaller role to certain shocks, relative to other shocks, in certain episodes. In contrast, shock-percentile restrictions assume that the realisation of a shock in some period was large relative to other realisations of the same shock, and thus do not directly constrain the role of other shocks in that period.

Shock-percentile restrictions can alternatively be interpreted as restrictions on the *contributions* of a specific shock to changes in a variable. For example, consider the shock-percentile restriction that

---

21 The event constraints proposed in LMN and Ludvigson *et al* (2017) can also be viewed as narrative restrictions, though they were proposed independently.

22 The restriction that a shock was the largest realisation of that shock over the sample period is nested within (9) when  $\alpha = 1$ . Restrictions on the absolute values of shocks can be implemented by replacing the shocks with their absolute values in (7) and (9).

the financial uncertainty shock in the Black Monday episode was larger than the 75th percentile of the historical distribution of financial uncertainty shocks. Given the sign normalisation on  $\mathbf{H}$ , this restriction is equivalent to imposing that the *contribution* of the financial uncertainty shock to the one-step-ahead forecast error in financial uncertainty in the Black Monday episode exceeds the 75th percentile of the historical distribution of these contributions. This follows from the fact that the rank ordering of the shocks is preserved when multiplying them by a positive constant (here, the impact impulse response of financial uncertainty to the shock, which is normalised to be positive). This interpretation would no longer hold in models where  $\mathbf{H}$  varies over time.

## 5. Uncertainty and Business Cycles Revisited

This section uses the shock-percentile restrictions to revisit the relationship between real activity and uncertainty in the United States, which is the empirical setting examined by LMN. Section 5.1 describes the collection of identifying restrictions imposed by LMN and shows that their shock-magnitude restrictions play a pivotal role in generating their key results. Section 5.2 re-examines these results when replacing the shock-magnitude restrictions with shock-percentile restrictions. Section 5.3 imposes an additional shock-percentile restriction related to the 1998 Russian Financial Crisis. Consistent with LMN, I focus on estimating identified sets and abstract from sampling uncertainty.<sup>23</sup>

### 5.1 Identifying restrictions in LMN

LMN impose two broad classes of identifying restrictions: 1) 'external variable constraints'; and 2) 'event constraints', which include shock-magnitude restrictions.

#### 5.1.1 External variable constraints

The external variable constraints require the structural shocks to be correlated with variables that are external to the VAR. These constraints do not constrain the external variables to be exogenous with respect to any shocks in the system, unlike much of the literature that makes use of 'external instruments' for identification (e.g. Mertens and Ravn 2013; Stock and Watson 2018).

Let  $\mathbf{S}_t = (S_{1t}, S_{2t})'$  denote a vector of external variables, where  $S_{1t}$  is a real stock market return and  $S_{2t}$  is the log difference in the real price of gold.<sup>24</sup> The external variable constraints require that

$$\text{cov}(S_{1t}, \boldsymbol{\varepsilon}_{jt}) \leq 0, \quad j = M, F \quad (13)$$

$$\text{cov}(S_{2t}, \boldsymbol{\varepsilon}_{jt}) \geq 0, \quad j = M, F. \quad (14)$$

That is, the two uncertainty shocks are required to be negatively correlated with the real stock market return and positively correlated with the change in the real price of gold. Since  $\boldsymbol{\varepsilon}_t = \mathbf{Q}'\boldsymbol{\Sigma}_{tr}^{-1}\mathbf{u}_t$ , it follows that  $\text{cov}(\mathbf{S}_t, \boldsymbol{\varepsilon}_t') = \text{cov}(\mathbf{S}_t, \mathbf{u}_t')(\boldsymbol{\Sigma}_{tr}^{-1})'\mathbf{Q}$ . The external variable constraints are therefore inequality restrictions on  $\mathbf{Q}$ . The restrictions depend on the covariances between the external variables  $\mathbf{S}_t$  and the reduced-form innovations  $\mathbf{u}_t$ ; these covariances constitute additional reduced-form parameters. In what follows, assume  $\boldsymbol{\phi}$  additionally contains  $\text{vec}(\text{cov}(\mathbf{S}_t, \mathbf{u}_t'))$ .

<sup>23</sup> LMN use a bootstrap to account for sampling uncertainty in one exercise (see their Figure 6), though the theoretical properties of their bootstrap are unknown. Section 6 presents an additional empirical application in which I account for sampling uncertainty using a prior-robust Bayesian method.

<sup>24</sup>  $S_{1t}$  is the Center for Research in Security Prices value-weighted stock market index return.  $S_{2t}$  is the log difference in the real price of gold (deflated using the US Consumer Price Index). See LMN for details about how these variables are constructed.

### 5.1.2 Event constraints

As discussed in Section 2.3, event constraints restrict the values of the structural shocks in specific periods. LMN impose the following event constraints:

- (i)  $\varepsilon_{F\tau_1}(\boldsymbol{\phi}, \mathbf{Q}, \mathbf{u}_{\tau_1}) \geq k_1$  at  $\tau_1 = 1987:M10$
- (ii)  $(\varepsilon_{F\tau_2}(\boldsymbol{\phi}, \mathbf{Q}, \mathbf{u}_{\tau_2}) \geq k_2) \vee (\varepsilon_{M\tau_2}(\boldsymbol{\phi}, \mathbf{Q}, \mathbf{u}_{\tau_2}) \geq k_3)$  at  $\tau_2 = 2008:M9$
- (iii)  $\varepsilon_{M\tau_3}(\boldsymbol{\phi}, \mathbf{Q}, \mathbf{u}_{\tau_3}) \geq k_4$  at  $\tau_3 = 1970:M12$
- (iv)  $\sum_{t \in \tau_4} \varepsilon_{Y_t}(\boldsymbol{\phi}, \mathbf{Q}, \mathbf{u}_t) \leq 0$  for  $\tau_4 \in \{2007:M12, \dots, 2009:M6\}$
- (v)  $\varepsilon_{M\tau_5}(\boldsymbol{\phi}, \mathbf{Q}, \mathbf{u}_{\tau_5}) \geq 0$  and  $\varepsilon_{F\tau_5}(\boldsymbol{\phi}, \mathbf{Q}, \mathbf{u}_{\tau_5}) \geq 0$  at  $\tau_5 = 1979:M10$
- (vi)  $\varepsilon_{M\tau_6}(\boldsymbol{\phi}, \mathbf{Q}, \mathbf{u}_{\tau_6}) \geq 0$  and  $\varepsilon_{F\tau_6}(\boldsymbol{\phi}, \mathbf{Q}, \mathbf{u}_{\tau_6}) \geq 0$  at  $\tau_6 \in \{2011:M7, 2011:M8\}$

Constraints (i)–(iii) are shock-magnitude restrictions. Constraint (i) requires the financial uncertainty shock in the Black Monday episode to exceed  $k_1$  standard deviations. Constraint (ii) requires the financial uncertainty shock and/or the macroeconomic uncertainty shock in September 2008 (the Lehman Brothers collapse) to exceed  $k_2$  and  $k_3$  standard deviations in size, respectively. Constraint (iii) requires the macroeconomic uncertainty shock in December 1970 (leading up to the collapse of the Bretton Woods system of fixed exchange rates) to exceed  $k_4$  standard deviations. LMN elicit the bounds  $\mathbf{k} = (k_1, \dots, k_4)'$  using the procedure described in Section 3. In the exercises below, I set these bounds equal to the 75th percentiles of the shock distributions implied by the RWZ algorithm, consistent with one of the specifications in LMN. In practice, this means that  $\mathbf{k} = (4.16, 4.57, 4.73, 4.05)'$ .

Constraints (iv)–(vi) are additional event constraints motivated by narrative information. Constraint (iv) requires the sum (equivalently, average) of the ‘real activity shocks’ from December 2007 to June 2009 to be positive, corresponding to the Great Recession. Constraints (v) and (vi) require both uncertainty shocks to be non-negative during October 1979 (the Volcker episode) and during the two months corresponding to the US debt ceiling crisis.

### 5.1.3 Numerical implementation

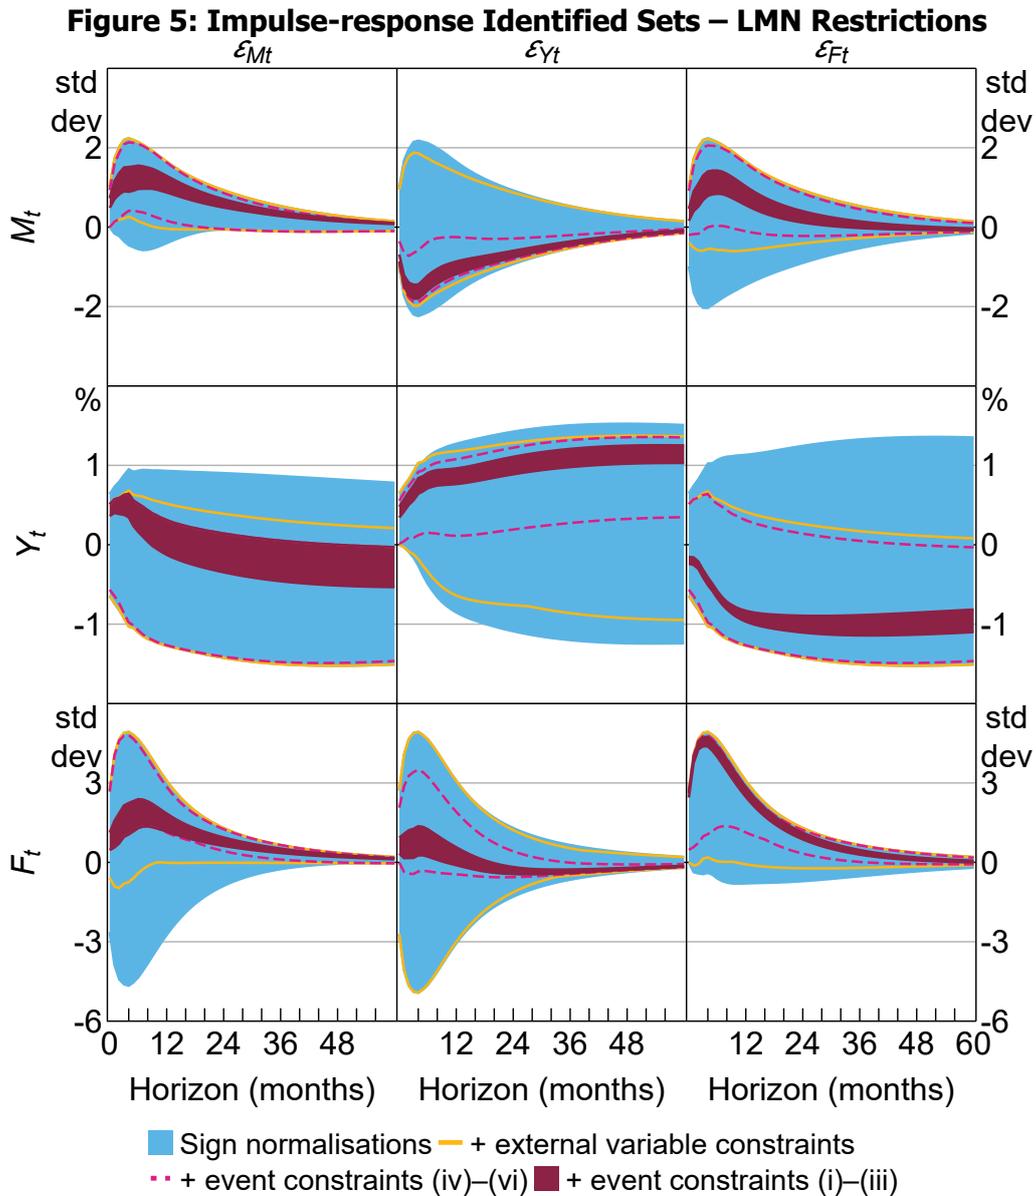
I approximate identified sets using a large number of draws of  $\mathbf{Q}$  from the uniform distribution over  $\mathcal{Q}(\boldsymbol{\phi}|S)$ .<sup>25</sup> Obtaining these draws using the RWZ algorithm and an accept-reject step (as is standard) is computationally burdensome, because  $\mathcal{Q}(\boldsymbol{\phi}|S)$  is ‘small’ relative to  $\mathcal{Q}(\boldsymbol{\phi})$ . I therefore employ the ‘soft sign restrictions’ approach developed in Read and Zhu (2025), which can be more computationally efficient when identification is tight. I use around 300,000 draws to approximate the identified sets, motivated by Montiel Olea and Nesbit (2021), who provide results about the number of draws required to approximate identified sets with a specified level of accuracy.<sup>26</sup> See Appendix C.1 for additional details.

<sup>25</sup> Under the identifying restrictions, it is not guaranteed that identified sets for the parameters of interest are convex. If the identified sets are non-convex, the interval estimates presented below can be interpreted as estimating convex hulls of identified sets. See Giacomini and Kitagawa (2021a, 2021b) for sufficient conditions under which impulse-response identified sets are convex.

<sup>26</sup> This contrasts with the treatment in LMN, who obtain 1.5 million draws of  $\mathbf{Q}$  from  $\mathcal{Q}(\boldsymbol{\phi})$  and approximate identified sets using only those draws that satisfy the identifying restrictions; for example, the identified sets under the event and external variable constraints described above are approximated using the 169 draws that satisfy the restrictions.

### 5.1.4 Results under LMN restrictions

Figure 5 presents identified sets for the impulse responses under different subsets of the identifying restrictions from LMN.<sup>27</sup> Under the full set of restrictions, a positive financial uncertainty shock unambiguously decreases output at all horizons. A positive macroeconomic uncertainty shock unambiguously increases output at short horizons, which is consistent with 'growth-options' theories of uncertainty. Positive shocks to real activity increase financial uncertainty at short horizons but lead to a persistent decrease in macroeconomic uncertainty. These results together suggest that elevated macroeconomic uncertainty during recessions is likely to reflect an endogenous response to adverse shocks associated with business cycle fluctuations. In contrast, heightened financial uncertainty may be an exogenous driver of recessions.



<sup>27</sup> The estimated identified sets are very similar to those presented in LMN, though they are slightly wider, reflecting the larger number of draws used to approximate identified sets and the consequent reduction in approximation error.

The results under weaker sets of restrictions make it clear that the shock-magnitude restrictions are crucial for obtaining the key results in LMN; in the absence of the shock-magnitude restrictions, identified sets for the output responses to both macroeconomic and financial uncertainty shocks span zero at all horizons. The credibility of the shock-magnitude restrictions is therefore a crucial ingredient in assessing the overall credibility of the results.

## 5.2 Shock-percentile restrictions

Consider replacing the shock-magnitude restrictions (i)–(iii) with shock-percentile restrictions.<sup>28</sup> More specifically, constraints (i)–(iii) are replaced with:

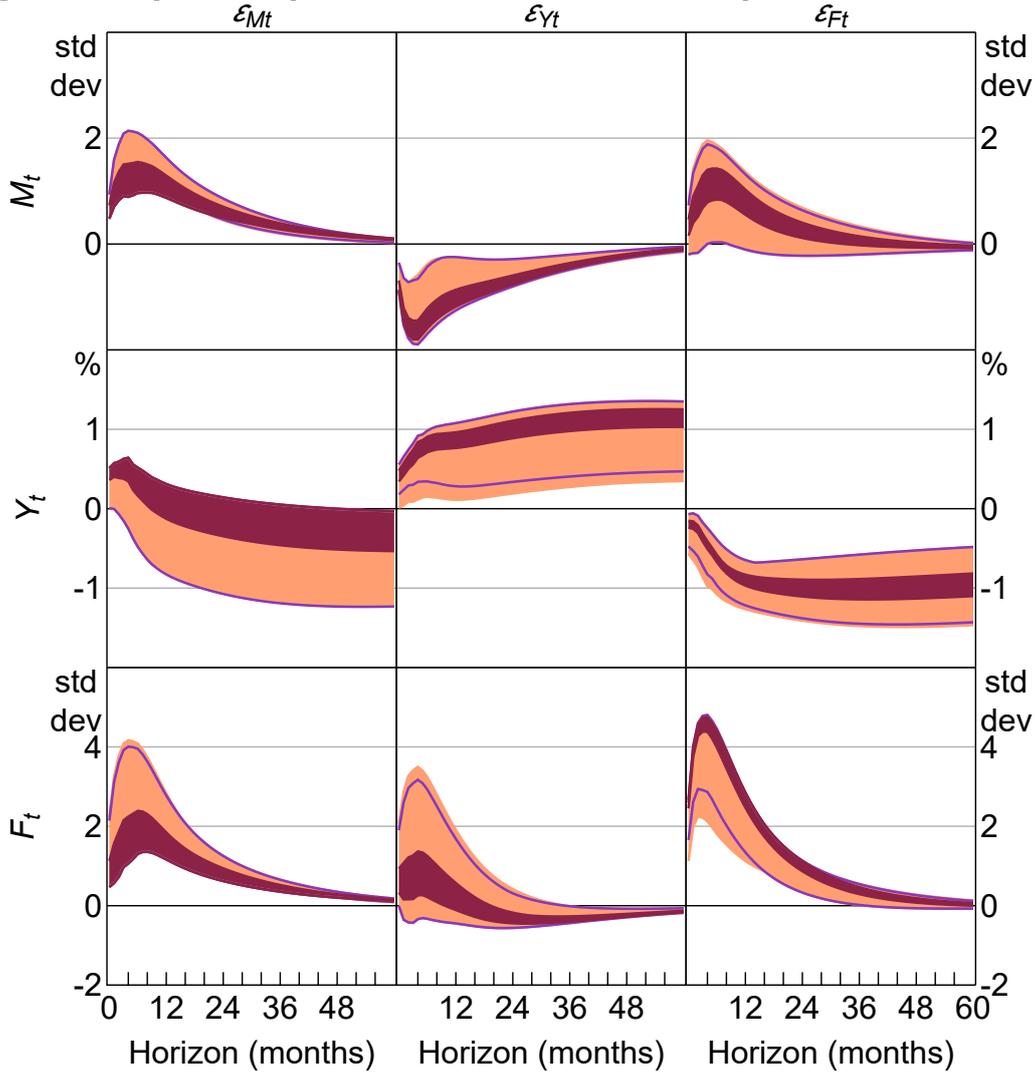
- (i\*)  $\varepsilon_{F\tau_1}(\phi, \mathbf{Q}, \mathbf{u}_{\tau_1}) \geq G_F(\alpha; \phi, \mathbf{Q}, \mathbf{U})$  at  $\tau_1 = 1987:M10$
- (ii\*)  $\left( \varepsilon_{F\tau_2}(\phi, \mathbf{Q}, \mathbf{u}_{\tau_2}) \geq G_F(\alpha; \phi, \mathbf{Q}, \mathbf{U}) \right) \vee \left( \varepsilon_{M\tau_2}(\phi, \mathbf{Q}, \mathbf{u}_{\tau_2}) \geq G_M(\alpha; \phi, \mathbf{Q}, \mathbf{U}) \right)$  at  $\tau_2 = 2008:M9$
- (iii\*)  $\varepsilon_{M\tau_3}(\phi, \mathbf{Q}, \mathbf{u}_{\tau_3}) \geq G_M(\alpha; \phi, \mathbf{Q}, \mathbf{U})$  at  $\tau_3 = 1970:M12$ .

I initially set  $\alpha = 0.75$ , so these shock-percentile restrictions are in a sense analogous to the shock-magnitude restrictions in Section 5.1.2. Constraint (i\*) requires the financial uncertainty shock in the Black Monday episode to exceed the 75th percentile of the shock's historical distribution. Constraint (ii\*) requires the financial uncertainty shock and/or the macroeconomic uncertainty shock in the Lehman episode to exceed the 75th percentiles of these shocks' respective historical distributions. Constraint (iii\*) requires the macroeconomic uncertainty shock in December 1970 to exceed the 75th percentile of the shock's historical distribution. These restrictions represent the belief that relatively large uncertainty shocks occurred in these periods, without taking a stand on the exact numeric bound that constitutes a 'large' shock.

### 5.2.1 Impulse responses

Figure 6 presents impulse-response identified sets when replacing shock-magnitude restrictions (i)–(iii) with shock-percentile restrictions (i\*)–(iii\*). The identified sets under the shock-magnitude restrictions are also presented for comparison. Under the shock-percentile restrictions, a positive financial uncertainty shock unambiguously leads to a decline in output at all horizons, which is consistent with the results in LMN. In contrast, identified sets for the output response to a macroeconomic uncertainty shock include zero at most horizons, and the sign of the output response is largely ambiguous. It remains the case that a positive output shock unambiguously decreases macroeconomic uncertainty, consistent with the idea that elevated macroeconomic uncertainty during recessions is likely to reflect an endogenous response to adverse shocks associated with business cycle fluctuations. In contrast, the sign of the response of financial uncertainty to a positive output shock is ambiguous at all horizons; in principle, heightened financial uncertainty during recessions could reflect the endogenous response of financial uncertainty to business cycle fluctuations.

<sup>28</sup>LMN also motivate the *timing* of their shock-magnitude restrictions based on the RWZ algorithm, which could be problematic when the identified timing of large shocks depends on the distribution for  $\mathbf{Q}$ . For example, the identified set for  $\arg \max_t \varepsilon_{Ft}(\phi, \mathbf{Q}, \mathbf{u}_t)$  consists of (at least) 13 different months, so the period where the largest financial uncertainty shock occurred is not pinned down by the data. However, the timing of the restrictions is additionally motivated using narrative information. I therefore proceed under the assumption that the timing of these restrictions is credible.

**Figure 6: Impulse-response Identified Sets – Shock-percentile Restrictions**

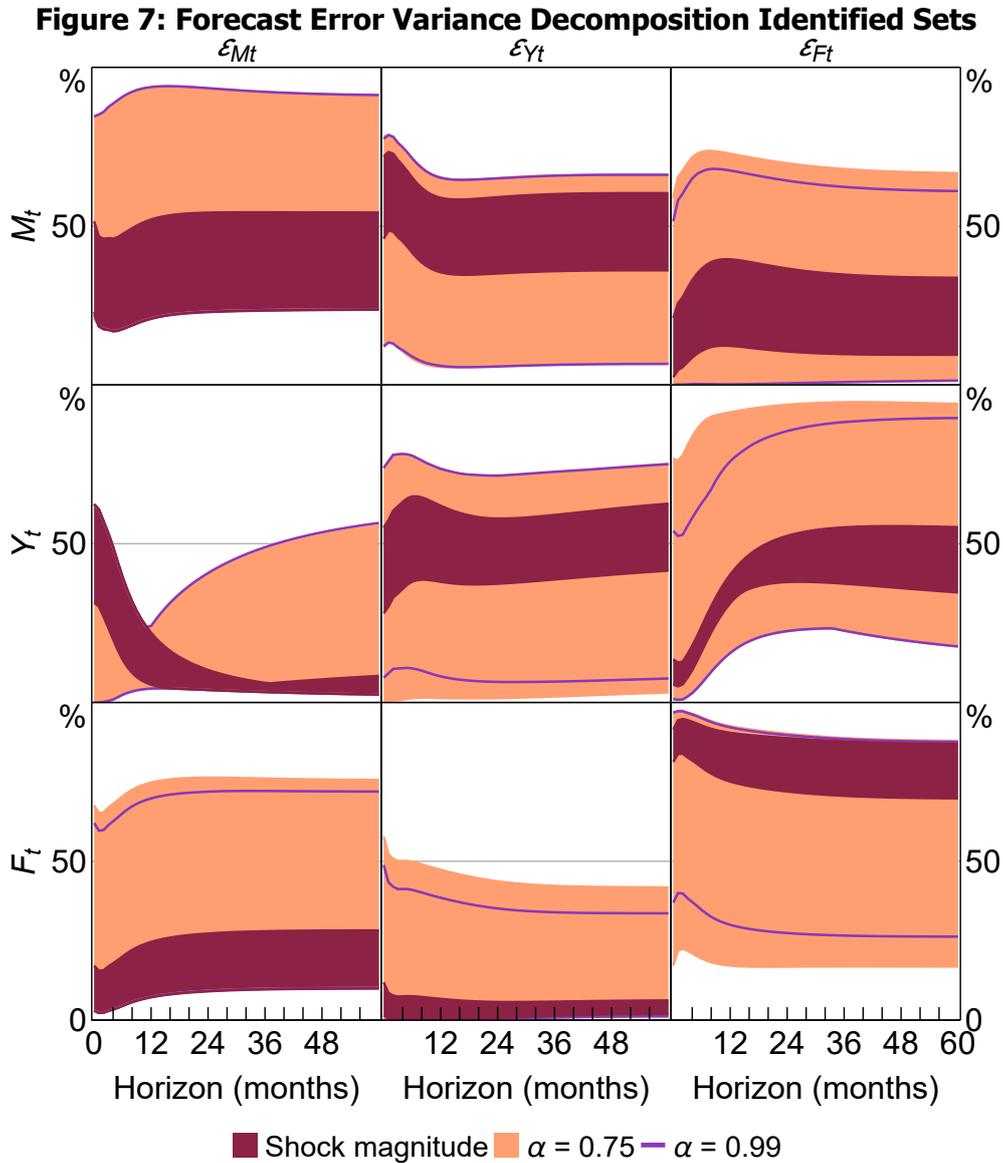
Notes: 'Shock magnitude' corresponds to identifying restrictions from LMN, with lower bounds on shock magnitudes based on 75th percentile of shock distribution from bound-elicitation procedure.  $\alpha = 0.75$  and  $\alpha = 0.99$  replace shock-magnitude restrictions (i)–(iii) with shock-percentile restrictions (i\*)–(iii\*).

To summarise, replacing the shock-magnitude restrictions with analogous shock-percentile restrictions yields qualitatively similar conclusions about the output effects of financial uncertainty shocks and the response of macroeconomic uncertainty to output shocks. In contrast, there is no longer unambiguous evidence about the output effects of macroeconomic uncertainty shocks or the response of financial uncertainty to output shocks.

Figure 6 also plots the responses obtained under stronger versions of the shock-percentile restrictions, where  $\alpha = 0.99$ . These restrictions require the financial and macroeconomic uncertainty shocks in the selected episodes to lie in the upper 1 per cent of the historical distributions of these shocks. Imposing these stronger restrictions narrows the identified sets a little in some cases, but overall there is little substantive difference relative to when  $\alpha = 0.75$ .

### 5.2.2 How important are uncertainty shocks?

To quantify the importance of different shocks for driving variation in output and uncertainty on average over time, Figure 7 presents identified sets for FEVDs. The identified sets are obtained under the shock-percentile restrictions with  $\alpha = 0.75$  and  $\alpha = 0.99$ . For comparison, the figure also presents identified sets based on the shock-magnitude restrictions.



Notes: 'Shock magnitude' corresponds to identifying restrictions from LMN, with lower bounds on shock magnitudes based on 75th percentile of shock distribution from bound-elicitation procedure.  $\alpha = 0.75$  and  $\alpha = 0.99$  replace shock-magnitude restrictions (i)–(iii) with shock-percentile restrictions (i\*)–(iii\*).

The shock-magnitude restrictions in LMN unambiguously imply that: 1) macroeconomic and financial uncertainty shocks are important drivers of output fluctuations at different horizons, with macroeconomic uncertainty shocks making substantial contributions at shorter horizons and financial uncertainty shocks at longer horizons; 2) the bulk of the variation in financial uncertainty is driven by exogenous shocks to financial uncertainty; and 3) a large share of the variation in macroeconomic uncertainty represents the endogenous response of macroeconomic uncertainty to other shocks.

Under the shock-percentile restrictions, the identified set for the contribution of macroeconomic uncertainty shocks to the forecast error variance of output includes values close to zero at all horizons. This means that there is no longer unambiguous evidence that macroeconomic uncertainty shocks are quantitatively important drivers of output fluctuations at short horizons. The identified sets now also admit the possibility that macroeconomic uncertainty shocks contribute materially to unexpected variation in output at longer horizons.

While the identified sets for the contribution of financial uncertainty shocks to output are substantially wider than obtained under the shock-magnitude restrictions, they still unambiguously imply that financial uncertainty shocks explain a non-trivial share of output fluctuations beyond short horizons; financial uncertainty shocks explain no less than about 20 per cent of unexpected changes in output at the 2–3 year horizon. So these shocks may indeed be an important driver of business cycles, consistent with the VAR analyses in Bloom (2009), Caggiano, Castelnuovo and Groshenny (2014), Leduc and Liu (2016) and Andreasen *et al* (2024), among others. The identified sets also admit the possibility that financial uncertainty shocks are not the *primary* driver of output fluctuations at business cycle frequencies, consistent with Brianti (2025).

In terms of what shocks drive measured uncertainty, the identified sets now admit the possibilities that financial uncertainty is largely driven by other (non-financial uncertainty) shocks and that macroeconomic uncertainty is primarily driven by macroeconomic uncertainty shocks. These qualitative conclusions hold both when  $\alpha = 0.75$  or when imposing stronger shock-percentile restrictions with  $\alpha = 0.99$ .<sup>29</sup>

Overall, relative to the shock-magnitude restrictions, the shock-percentile restrictions deliver less-conclusive evidence about the importance of uncertainty shocks for driving measured uncertainty and real activity.

### 5.2.3 Understanding the results

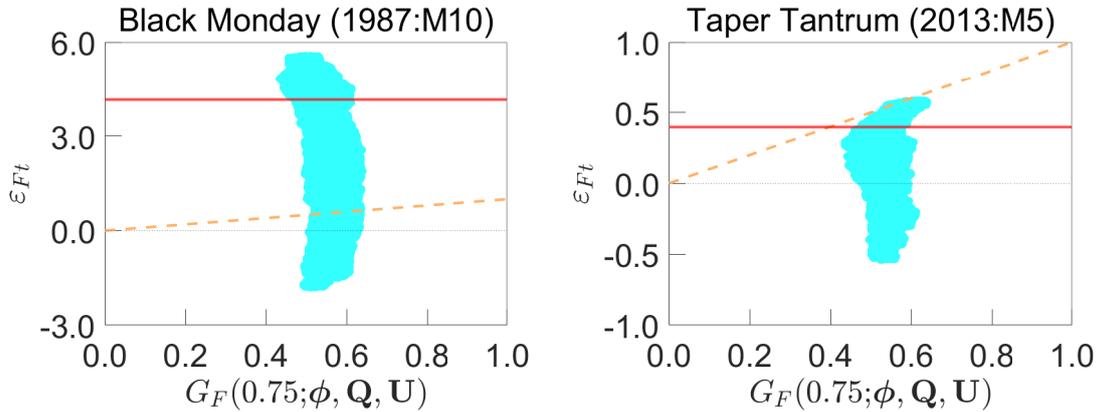
While the shock-percentile restrictions in this example generate wider identified sets than the shock-magnitude restrictions, the results are obtained without relying on the arbitrary algorithm used to simulate parameter values. In that sense, replacing shock-magnitude restrictions with shock-percentile restrictions in this application can be viewed as an example of the trade-off between ‘credibility’ and ‘certitude’ discussed in Manski (2003, 2011). But why do the shock-percentile restrictions generate wider identified sets than the shock-magnitude restrictions?

To examine this question, the left panel of Figure 8 visualises the *joint* unconstrained identified set for the financial uncertainty shock in the Black Monday episode ( $\varepsilon_{F\tau_1}(\phi, \mathbf{Q}, \mathbf{u}_{\tau_1})$ ) and the 75th percentile of the historical distribution of financial uncertainty shocks ( $G_F(0.75; \phi, \mathbf{Q}, \mathbf{U})$ ). Each point in the shaded region represents a pair  $(\varepsilon_{F\tau_1}(\phi, \mathbf{Q}, \mathbf{u}_{\tau_1}), G_F(0.75; \phi, \mathbf{Q}, \mathbf{U}))$  corresponding to a value of  $\mathbf{Q}$  in the unconstrained identified set  $\mathcal{Q}(\phi)$ . Points lying above the dashed line satisfy the shock-percentile restriction  $\varepsilon_{F\tau_1}(\phi, \mathbf{Q}, \mathbf{u}_t) \geq G_F(0.75; \phi, \mathbf{Q}, \mathbf{U})$ . Points lying above the solid horizontal line satisfy the shock-magnitude restriction  $\varepsilon_{F\tau_1}(\phi, \mathbf{Q}, \mathbf{u}_t) \geq k_1$ . It is evident that there are no values of

<sup>29</sup> Appendices C.2 and C.3 apply the shock-percentile restrictions to revisit additional empirical questions considered in LMN. Appendix C.2 reaffirms the results in LMN that the data are inconsistent with recursive identification schemes, which have been widely used to identify uncertainty shocks. Appendix C.3 presents evidence that the structural shocks possess non-Gaussian features, which is inconsistent with the common Gaussianity assumption in Bayesian SVARs.

$\mathbf{Q} \in \mathcal{Q}(\phi)$  where the shock percentile exceeds the bound on the shock; in this sense, the shock-percentile restriction is strictly weaker than the shock-magnitude restriction. This is also the case for the other periods where these restrictions are imposed. Hence, the shock-percentile restrictions, when combined with the other identifying restrictions, generate wider identified sets than the shock-magnitude restrictions.

**Figure 8: Joint Unconstrained Identified Sets for Financial Uncertainty Shocks and Percentiles**



Notes: Shaded region represents joint unconstrained identified set for financial uncertainty shock ( $\epsilon_{Ft}$ ) in selected episode and historical shock percentiles ( $G_F(0.75; \phi, \mathbf{Q}, \mathbf{U})$ ). Points lying above the dashed line satisfy the shock-percentile restriction and points lying above the solid line satisfy the shock-magnitude restriction.

A natural question is: will it always be the case that a shock-percentile restriction is weaker than an analogous shock-magnitude restriction when the bound is elicited using the procedure in LMN? The answer is no. To illustrate, consider the financial uncertainty shock in the Taper Tantrum episode ( $\tau_7 = 2013:M5$ ); this was a period of volatility in financial markets, which was triggered by suggestions that the Federal Reserve might reduce the pace of bond purchases under its quantitative easing program. The right panel of Figure 8 shows that all values of  $\mathbf{Q} \in \mathcal{Q}(\phi)$  that satisfy a shock-percentile restriction based on this episode also satisfy the analogous shock-magnitude restriction. The shock-percentile restriction would therefore be unambiguously stronger than the shock-magnitude restriction.

### 5.3 Additional restrictions

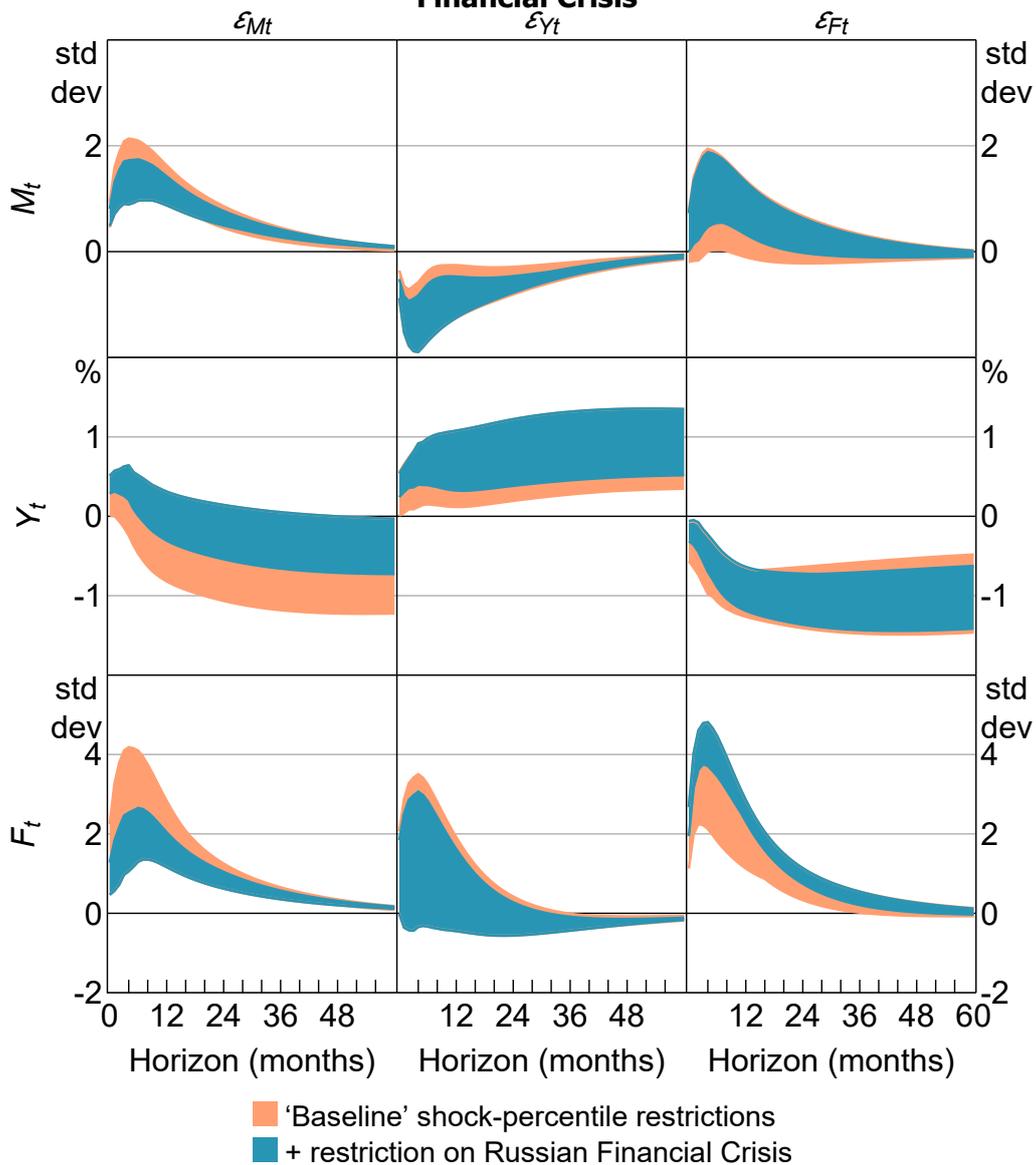
The shock-percentile restrictions considered above yield conclusions about the effects of uncertainty shocks that are qualitatively consistent, along some dimensions, with the results obtained under the shock-magnitude restrictions. However, as discussed, the identified sets tend to be substantially wider, so while the results are arguably more credible, they are in a sense less economically informative. It is therefore valuable to impose additional credible identifying restrictions to sharpen identification.

To that end, I impose an additional shock-percentile restriction related to the 1998 Russian Financial Crisis. More specifically, I restrict the financial uncertainty shock in August 1998 to exceed the 75th percentile of the historical distribution of uncertainty shocks. This is the month in which Russia defaulted on ruble-denominated government debt and devalued the ruble, triggering a spike in uncertainty in financial markets; for example, the CBOE Volatility Index (VIX) roughly doubled over August. This episode led to the near collapse of the hedge fund Long-Term Capital Management,

which was subsequently (in September 1998) bailed out by a consortium of private banks in a deal facilitated by the Federal Reserve (e.g. Lowenstein 2000).

Augmenting the restrictions from Section 5.2 with the additional shock-percentile restriction substantially sharpens identification of the impulse responses to the uncertainty shocks, better disentangling shocks to macroeconomic and financial uncertainty (Figure 9). For example, output unambiguously increases at short horizons following a positive macroeconomic uncertainty shock, and macroeconomic uncertainty unambiguously increases following a positive financial uncertainty shock; these results are broadly consistent with the results from LMN.

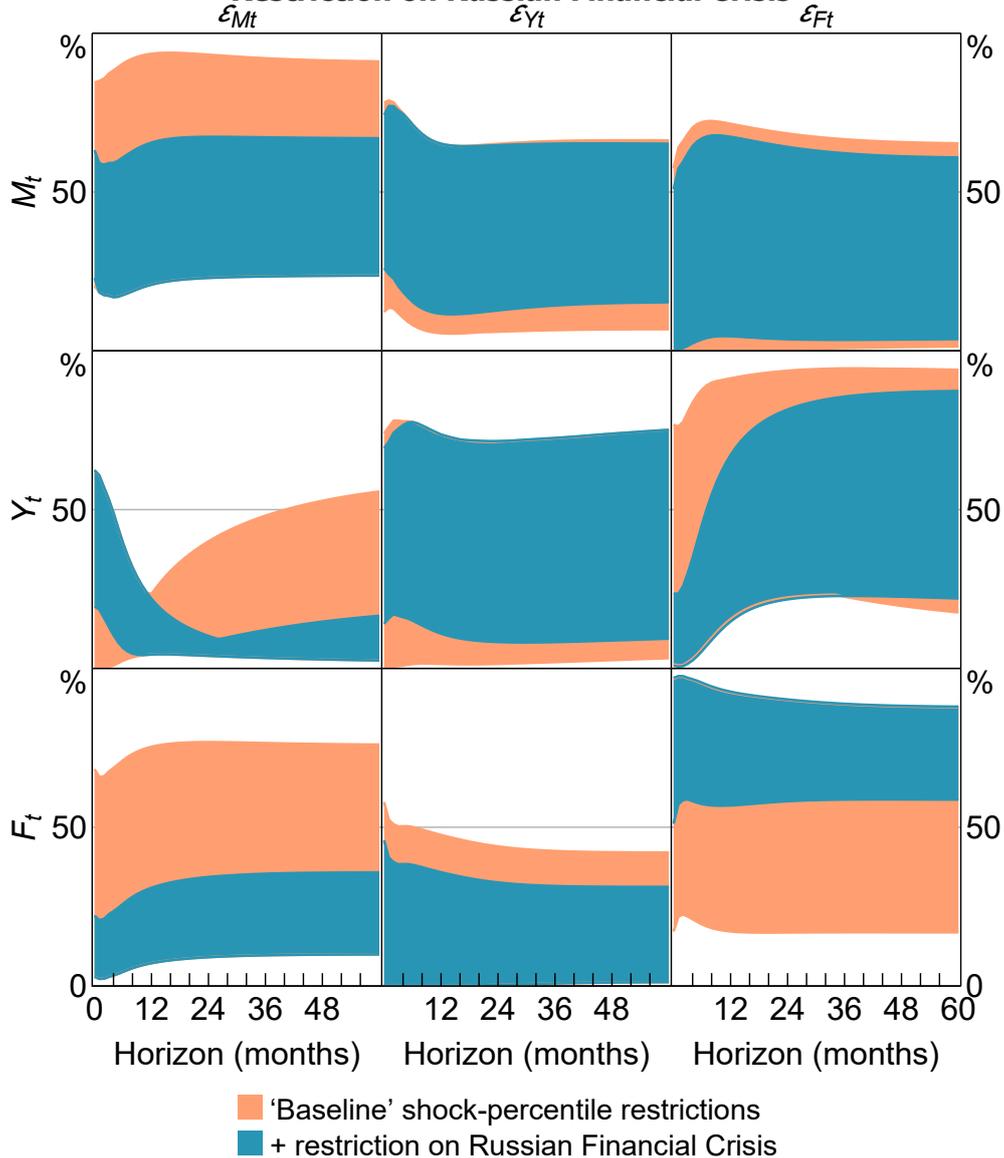
**Figure 9: Impulse-response Identified Sets – Shock-percentile Restriction on Russian Financial Crisis**



The additional shock-percentile restriction also tightens identified sets for the FEVDs, delivering less-ambiguous evidence about the importance of the different shocks for driving variation in output and measured uncertainty (Figure 10). More specifically: financial uncertainty shocks unambiguously explain at least half of the forecast error variance of financial uncertainty at all horizons, so most of the variation in financial uncertainty is driven by exogenous shocks to financial uncertainty; a

substantial share of the variation in macroeconomic uncertainty is now unambiguously ascribed to other shocks; and macroeconomic uncertainty shocks unambiguously make a material contribution to output fluctuations at short horizons and a limited contribution at longer horizons, and vice versa for financial uncertainty shocks. These results are broadly consistent with those obtained under LMN's shock-magnitude restrictions. Importantly, however, they are obtained without specifying an arbitrary numeric lower bound on shock magnitudes, and so may be viewed as more credible.

**Figure 10: Forecast Error Variance Decomposition Identified Sets – Shock-percentile Restriction on Russian Financial Crisis**



## 6. Estimating the Effects of US Monetary Policy

This section applies a shock-percentile restriction to help identify the effects of US monetary policy, building on Antolín-Díaz and Rubio-Ramírez (2018) (henceforth, AR18). They identify US monetary policy shocks using a combination of sign restrictions on impulse responses (as in Uhlig (2005)) and narrative restrictions. I consider augmenting their narrative restrictions with a shock-percentile restriction, reflecting the belief that the monetary policy shock in October 1979 (the Volcker episode)

was large relative to the historical distribution of monetary policy shocks, consistent with the view that the shock in this episode represented a ‘major anti-inflationary shock to monetary policy’ (Romer and Romer 1989, p 142).

## 6.1 Reduced-form VAR

The reduced-form VAR is from Uhlig (2005). The endogenous variables are real GDP, the GDP deflator, a commodity price index, total reserves, non-borrowed reserves (all in natural logarithms) and the federal funds rate. The data are monthly, with the sample beginning in January 1965 and ending in November 2007. The VAR includes 12 lags. As discussed further below, I estimate the model using a prior-robust Bayesian method under the assumption that the structural shocks are normally distributed. The prior for the reduced-form parameters is the uninformative Jeffreys’ prior, so the posterior is normal-inverse-Wishart.

## 6.2 Identifying restrictions

As identifying restrictions, AR18 impose a mix of sign restrictions on impulse responses and narrative restrictions. The sign restrictions follow Uhlig (2005): the response of the federal funds rate to a positive monetary policy shock is non-negative for  $h = 0, 1, \dots, 5$ ; and the responses of the GDP deflator, the commodity price index and non-borrowed reserves are non-positive for  $h = 0, 1, \dots, 5$ . As narrative restrictions, they impose that the monetary policy shock in October 1979 was positive and was the ‘overwhelming contributor’ to the observed unexpected change in the federal funds rate. The narrative restrictions represent the belief that the monetary policy shock in this episode was large relative to other shocks in the same period in terms of the shocks’ contributions to the change in the federal funds rate. In an additional exercise involving a richer set of narrative restrictions, the restrictions are that the monetary policy shock was: positive in April 1974, October 1979, December 1988 and February 1994; negative in December 1990, October 1998, April 2001 and November 2002; and the ‘most important contributor’ to the observed unexpected change in the federal funds rate in these months.<sup>30</sup>

To examine the identifying power of a shock-percentile restriction in this setting, I additionally impose the shock-percentile restriction that the monetary policy shock in the Volcker episode exceeded the 90th percentile of the historical distribution of monetary policy shocks.<sup>31</sup> Unlike the restrictions on the historical decomposition, this shock-percentile restriction represents the belief that the monetary policy shock in the Volcker episode was large relative to monetary policy shocks occurring in other periods. Imposing both restrictions in this episode therefore represents the belief that the monetary policy shock was ‘large’ along two dimensions: 1) its contribution to the change in the federal funds rate in this period relative to the contributions of other shocks; and 2) relative to other periods.

---

30 Let  $H_{i,j,t}$  be the contribution of shock  $j$  to the one-step-ahead forecast error in variable  $i$  in period  $t$ . The restriction that the first shock was the ‘overwhelming contributor’ to variable  $i$  in period  $t$  is  $|H_{i,1,t}| \geq \sum_{j \neq 1} |H_{i,j,t}|$  and the restriction that it was the ‘most important contributor’ to variable  $i$  in period  $t$  is  $|H_{i,1,t}| \geq \max_{j \neq 1} |H_{i,j,t}|$ .

31 Giacomini *et al* (2021) and Read (2022) impose the ‘shock-ranking’ restriction that the monetary policy shock in the Volcker episode was the largest absolute realisation of the shock in the sample period, which is stronger than the shock-percentile restriction considered here.

### 6.3 Results

I estimate the impulse responses to a 100 basis point monetary policy shock both under the identifying restrictions from AR18 and when additionally imposing the shock-percentile restriction. To do this, I use the prior-robust Bayesian approach to inference developed in Giacomini and Kitagawa (2021a) and extended to the case of narrative restrictions in Giacomini *et al.* (2023). This approach to inference provides a tractable way to account for sampling uncertainty when estimating identified sets; the 'set of posterior medians' under this approach can be interpreted as an estimator of the identified set and the 'robust credible interval' can be interpreted as an asymptotically valid frequentist confidence interval for the identified set.<sup>32</sup>

Figure 11 plots the impulse responses of the federal funds rate and output to a monetary policy shock that increases the federal funds rate by 100 basis points on impact. Under the original Volcker restrictions (left panels), the set of posterior medians for the output response contains zero at almost all horizons considered.<sup>33</sup> In contrast, when augmenting these restrictions with the shock-percentile restriction, the set of posterior medians excludes zero at most horizons within the first three years. One way to quantify the evidence about the output response is to use the 'posterior lower probability'; this is the smallest posterior probability assigned to a particular hypothesis under the 'class of posteriors' underlying the robust Bayesian approach to inference. The posterior lower probability that the output response is negative at the six-month horizon is less than 1 per cent under the original restrictions, whereas it is around 65 per cent when additionally imposing the shock-percentile restriction. Hence, there is much stronger evidence of a decline in output at this (relatively short) horizon following a monetary policy shock. Overall, augmenting the baseline narrative restrictions with the shock-percentile restriction yields stronger evidence that output falls following a positive monetary policy shock, particularly at shorter horizons.

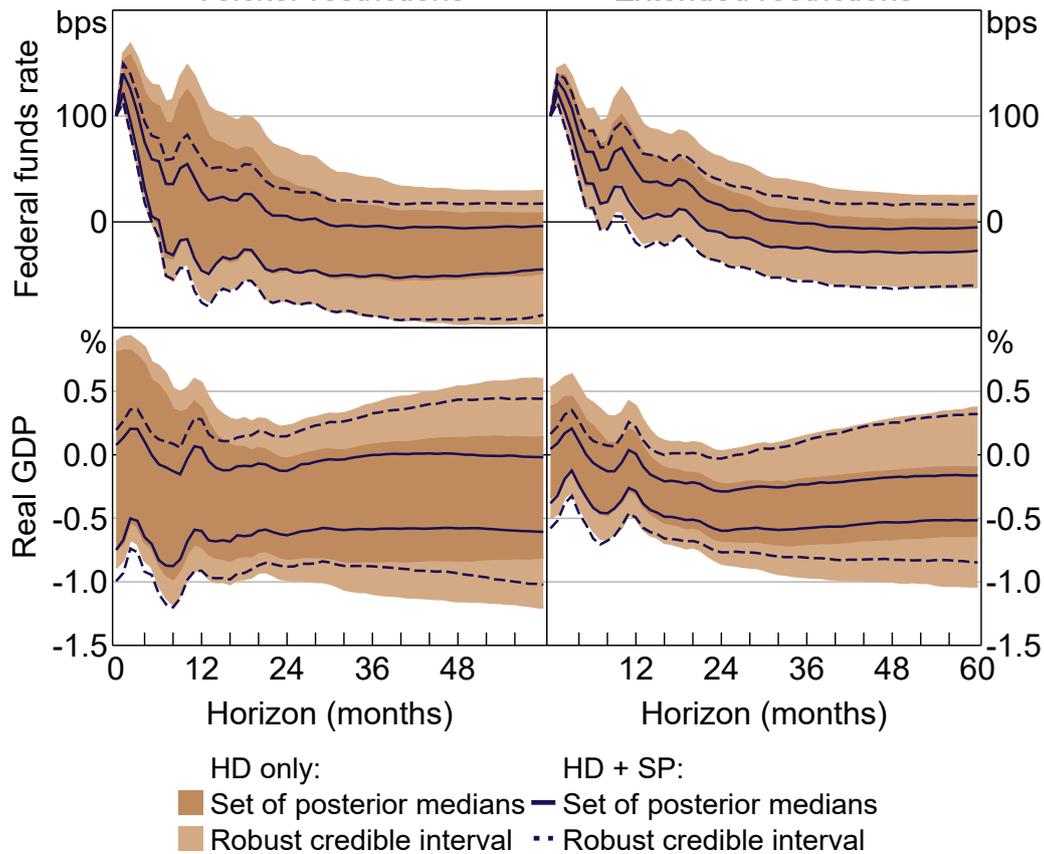
Under the extended set of restrictions from AR18 (right panels, Figure 11), the set of posterior medians includes zero at all horizons within the first year or so. This is no longer the case when additionally imposing the shock-percentile restriction; the estimator of the identified set unambiguously points to an output decline at almost all horizons beyond the first six months. The posterior lower probability that the output response is negative at the six-month horizon is about 65 per cent when imposing the shock-percentile restriction, compared with only 14 per cent under the original set of restrictions. So, again, there is much stronger evidence of a decline in output at short horizons when additionally imposing the shock-percentile restriction.

---

<sup>32</sup> See Appendix C.1 for details about how the approach to inference is implemented numerically.

<sup>33</sup> The estimates under the Volcker restrictions from AR18 coincide with those presented in the Online Appendix to Read (forthcoming). The results presented here are not directly comparable to those in AR18, because the impulse responses are normalised differently.

**Figure 11: Impulse Responses to Monetary Policy Shock**  
 Volcker restrictions      Extended restrictions



Notes: Impulse responses are to a 100 basis point monetary policy shock. 'HD only' represents results under identifying restrictions from Antolín-Díaz and Rubio-Ramírez (2018); 'HD + SP' additionally imposes shock-percentile restriction. Robust credible intervals are at the 68 per cent level.

To summarise, imposing a single shock-percentile restriction related to the Volcker episode substantially sharpens identification of the output effects of US monetary policy, yielding stronger evidence that output declines following a positive monetary policy shock, particularly at short horizons. This is the case even when imposing this restriction on top of an already rich set of narrative restrictions.

## 7. Conclusion

Identifying SVARs by imposing numeric bounds on shocks is inherently difficult. Instead, shock-percentile restrictions are based on beliefs about the size of shocks in selected episodes relative to the shock's historical distribution. These restrictions may therefore provide a more credible approach to sharpening identification in settings where narrative information suggests that shocks in selected episodes were relatively large, but the precise magnitudes of these shocks are uncertain.

## Appendix A: Analytical Results and Illustration in Bivariate Model

Section 3 argues that the bound-elicitation procedure used in formulating shock-magnitude restrictions generates arbitrary bounds, which casts doubt on the credibility of the restrictions. This appendix provides additional details to support this argument. Section A.1 introduces a simple bivariate model and uses it to show analytically that the unconstrained identified sets for the structural shocks always include zero. This result is then extended to a general  $n$ -variable SVAR. Section A.2 uses the bivariate model to illustrate the influence of the algorithm used to randomly draw parameter values in the bound-elicitation procedure.

### A.1 Unconstrained identified sets for shocks

This section proceeds by first showing that unconstrained identified sets for structural shocks always include zero using the simplest possible SVAR – a bivariate SVAR with no dynamics – as an example, since the unconstrained identified set can be characterised analytically in this case. I then extend this result to the case of a general  $n$ -variable SVAR.

Consider the bivariate SVAR(0),  $\mathbf{y}_t = \mathbf{H}\boldsymbol{\varepsilon}_t$ , where  $\mathbf{y}_t = (y_{1t}, y_{2t})'$  and  $\boldsymbol{\varepsilon}_t = (\varepsilon_{1t}, \varepsilon_{2t})'$ . Reparameterise  $\mathbf{H} = \boldsymbol{\Sigma}_{tr}\mathbf{Q}$ , where

$$\boldsymbol{\Sigma}_{tr} = \begin{bmatrix} \sigma_{11} & 0 \\ \sigma_{21} & \sigma_{22} \end{bmatrix} \quad (\sigma_{11}, \sigma_{22} > 0) \quad (\text{A1})$$

and  $\mathbf{Q}$  is an orthonormal matrix in the space of  $2 \times 2$  orthonormal matrices,  $\mathcal{O}(2)$ .  $\mathbf{Q}$  can be represented as

$$\mathbf{Q} \in \mathcal{O}(2) = \left\{ \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \right\} \cup \left\{ \begin{bmatrix} \cos \theta & \sin \theta \\ \sin \theta & -\cos \theta \end{bmatrix} \right\} \quad (\text{A2})$$

where  $\theta \in [-\pi, \pi]$ . This formulation of the model, which follows Baumeister and Hamilton (2015), means that the structural parameters can be expressed as functions of  $\boldsymbol{\Sigma}_{tr}$  and  $\theta$ . Restrictions on the structural parameters and/or the structural shocks are then restrictions on  $\theta$ .

In the absence of identifying restrictions, the identified set for  $\mathbf{H}$  is

$$\mathbf{H} \in \left\{ \begin{bmatrix} \sigma_{11} \cos \theta & -\sigma_{11} \sin \theta \\ \sigma_{21} \cos \theta + \sigma_{22} \sin \theta & \sigma_{22} \cos \theta - \sigma_{21} \sin \theta \end{bmatrix} \right\} \cup \left\{ \begin{bmatrix} \sigma_{11} \cos \theta & \sigma_{11} \sin \theta \\ \sigma_{21} \cos \theta + \sigma_{22} \sin \theta & \sigma_{21} \sin \theta - \sigma_{22} \cos \theta \end{bmatrix} \right\} \quad (\text{A3})$$

and the identified set for  $\mathbf{H}^{-1}$  is

$$\mathbf{H}^{-1} \in \left\{ \frac{1}{\sigma_{11}\sigma_{22}} \begin{bmatrix} \sigma_{22} \cos \theta - \sigma_{21} \sin \theta & \sigma_{11} \sin \theta \\ -\sigma_{21} \cos \theta - \sigma_{22} \sin \theta & \sigma_{11} \cos \theta \end{bmatrix} \right\} \cup \left\{ \frac{1}{\sigma_{11}\sigma_{22}} \begin{bmatrix} \sigma_{22} \cos \theta - \sigma_{21} \sin \theta & \sigma_{11} \sin \theta \\ \sigma_{22} \sin \theta + \sigma_{21} \cos \theta & -\sigma_{11} \cos \theta \end{bmatrix} \right\} \quad (\text{A4})$$

The sign normalisation  $\text{diag}(\mathbf{H}) \geq \mathbf{0}_{2 \times 1}$  generates an unconstrained identified set for  $\theta$ ,  $IS_\theta(\boldsymbol{\phi}) = [-\pi/2, \pi/2]$ .<sup>34</sup>

Given a value of the parameters,  $\boldsymbol{\phi}$  and  $\theta$ , and the realisation of the data in period  $t$ ,  $\mathbf{y}_t$ , we can use  $\varepsilon_{1t} = \mathbf{e}'_{1,2} \mathbf{H}^{-1} \mathbf{y}_t$  and the expression for  $\mathbf{H}^{-1}$  in (A4) to write  $\varepsilon_{1t}$  as

$$\varepsilon_{1t}(\boldsymbol{\phi}, \theta, \mathbf{y}_t) = (\sigma_{11} \sigma_{22})^{-1} (\sigma_{22} y_{1t} \cos \theta + (\sigma_{11} y_{2t} - \sigma_{21} y_{1t}) \sin \theta) \quad (\text{A5})$$

$IS_\theta(\boldsymbol{\phi})$  induces a conditional identified set for  $\varepsilon_{1t}$ :  $CIS_{\varepsilon_{1t}}(\boldsymbol{\phi}, \mathbf{y}_t) = \{\varepsilon_{1t}(\boldsymbol{\phi}, \theta, \mathbf{y}_t) : \theta \in IS_\theta(\boldsymbol{\phi})\}$ . Evaluating  $\varepsilon_{1t}(\boldsymbol{\phi}, \theta, \mathbf{y}_t)$  at the end points of  $IS_\theta(\boldsymbol{\phi})$  yields

$$\varepsilon_{1t}(\boldsymbol{\phi}, -\pi/2, \mathbf{y}_t) = -\frac{y_{1t}}{\sigma_{11}} C(\boldsymbol{\phi}, \mathbf{y}_t) \quad (\text{A6})$$

$$\varepsilon_{1t}(\boldsymbol{\phi}, \pi/2, \mathbf{y}_t) = \frac{y_{1t}}{\sigma_{11}} C(\boldsymbol{\phi}, \mathbf{y}_t) \quad (\text{A7})$$

Since  $\varepsilon_{1t}(\boldsymbol{\phi}, -\pi/2, \mathbf{y}_t) = -\varepsilon_{1t}(\boldsymbol{\phi}, \pi/2, \mathbf{y}_t)$ , there exist two values of  $\varepsilon_{1t}$  with opposite signs within the conditional identified set. It immediately follows that  $CIS_{\varepsilon_{1t}}(\boldsymbol{\phi}, \mathbf{y}_t)$  contains zero at any value of the reduced-form parameters and at any realisation of the data.<sup>35</sup>

Next, consider an  $n$ -variable SVAR. Given an arbitrary value of  $\boldsymbol{\phi}$  and  $\mathbf{u}_t$ , the conditional identified set for  $\varepsilon_{1t}$  includes zero if and only if the following set is non-empty:

$$\{\mathbf{q}_1 \in \mathbb{R}^n : \mathbf{e}'_{1,n} \boldsymbol{\Sigma}_{tr} \mathbf{q}_1 \geq 0, (\boldsymbol{\Sigma}_{tr}^{-1} \mathbf{u}_t)' \mathbf{q}_1 = 0\} \quad (\text{A8})$$

where the first inequality corresponds to the sign normalisation.<sup>36</sup> Let

$$\mathbf{W} = \begin{bmatrix} \mathbf{e}'_{1,n} \boldsymbol{\Sigma}_{tr} \\ (\boldsymbol{\Sigma}_{tr}^{-1} \mathbf{u}_t)' \end{bmatrix} \quad (\text{A9})$$

If  $n > 2$ , we can always find  $\mathbf{q}_1 \in \mathbb{R}^n$  satisfying  $\mathbf{W} \mathbf{q}_1 = \mathbf{0}_{2 \times 1}$  by computing an orthonormal basis for the null space of  $\mathbf{W}$ ; by the rank-nullity theorem, this null space has dimension  $n - \text{rank}(\mathbf{W}) \geq 1$ . And if some  $\mathbf{q}_1$  satisfies  $\mathbf{W} \mathbf{q}_1 = \mathbf{0}_{2 \times 1}$  then it clearly also satisfies  $\mathbf{W} \mathbf{q}_1 \geq \mathbf{0}_{2 \times 1}$ . The case where  $n = 2$  and  $\text{rank}(\mathbf{W}) \leq 2$  corresponds to the bivariate case above, in which it has already been shown that the identified set for the shock includes zero.  $\boldsymbol{\Sigma}_{tr}$  and  $\mathbf{u}_t$  are arbitrary, so the unconstrained identified set for  $\varepsilon_{1t}$  includes zero at any values of the reduced-form parameters and data.<sup>37</sup>

34 The unconstrained identified set for  $\theta$  is

$$\{\theta : \sigma_{11} \cos \theta \geq 0, \sigma_{22} \cos \theta \geq \sigma_{21} \sin \theta\} \cup \{\theta : \sigma_{11} \cos \theta \geq 0, \sigma_{21} \sin \theta \geq \sigma_{22} \cos \theta\}$$

which simplifies to  $\{\theta : \cos \theta \geq 0\}$ .

35 An alternative (more tedious) way to show this is to explicitly characterise the conditional identified set. Since  $IS_\theta(\boldsymbol{\phi})$  is an interval and  $\varepsilon_{1t}(\boldsymbol{\phi}, \theta, \mathbf{y}_t)$  is a continuous function of  $\theta$ ,  $CIS_{\varepsilon_{1t}}(\boldsymbol{\phi}, \mathbf{y}_t)$  is also an interval, so it can be characterised by its end points. The bounds of  $CIS_{\varepsilon_{1t}}(\boldsymbol{\phi}, \mathbf{y}_t)$  occur either at the end points of  $IS_\theta(\boldsymbol{\phi})$  (i.e.  $-\pi/2$  or  $\pi/2$ ) or at a critical point of  $\varepsilon_{1t}(\boldsymbol{\phi}, \theta, \mathbf{y}_t)$  in the interior of  $IS_\theta(\boldsymbol{\phi})$ . We could therefore characterise the bounds of  $CIS_{\varepsilon_{1t}}(\boldsymbol{\phi}, \mathbf{y}_t)$  by comparing the values of  $\varepsilon_{1t}(\boldsymbol{\phi}, \theta, \mathbf{y}_t)$  evaluated at the critical point and the end points.

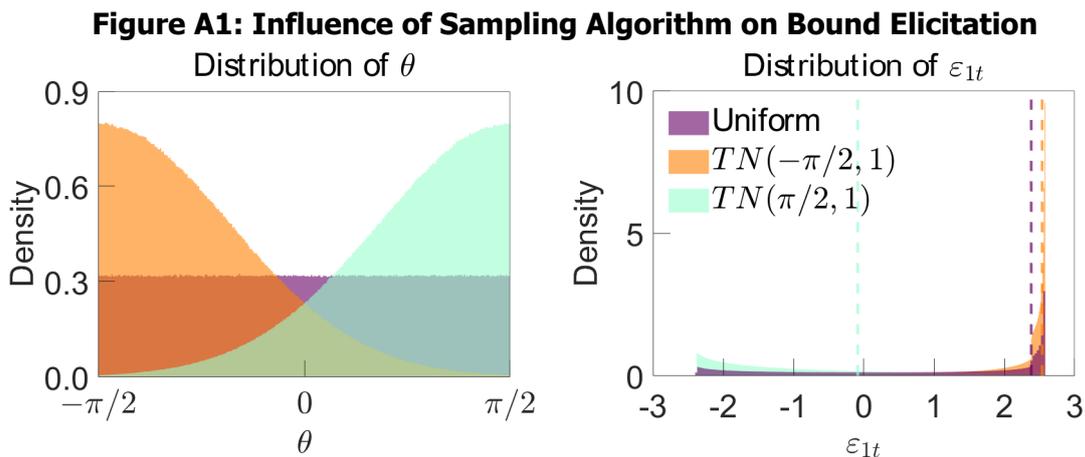
36 The sign normalisations on the remaining columns of  $\mathbf{Q}$  do not constrain  $\mathbf{q}_1$ ; given any  $\mathbf{q}_1 \in \mathbb{R}^n$ , we can always iteratively construct the remaining columns of  $\mathbf{Q}$  such that they are orthogonal to  $\mathbf{q}_1$  (and each other) and satisfy the sign normalisations.

37 This result closely parallels Proposition 3.2 in Read (forthcoming), who derives a sufficient condition under which impulse-response identified sets include zero when there are sign and zero restrictions on a single column of  $\mathbf{Q}$ .

## A.2 Influence of RWZ algorithm on bound elicitation

Section 3 argues that bounds on shock magnitudes elicited using the procedure in LMN could in principle be sensitive to the algorithm used to randomly generate parameter values. This is problematic, because the choice of algorithm is arbitrary in this setting. Using the bivariate model introduced in the previous section, this section illustrates the sensitivity of the bound-elicitation procedure to the choice of algorithm.

In the bivariate model, the distribution for  $\mathbf{Q}$  induced by the RWZ algorithm corresponds to a uniform distribution for  $\theta$  over its identified set; that is,  $\theta \sim U(-\pi/2, \pi/2)$  (e.g. Baumeister and Hamilton 2015). Given a realisation of the data  $\mathbf{y}_t$ , the distribution for  $\theta$  induces a distribution for  $\varepsilon_{1t}(\boldsymbol{\phi}, \theta, \mathbf{y}_t)$  via the transformation in (A5). Figure A1 plots these distributions given the same DGP used in Section 4.2 and a single realisation of the data.<sup>38</sup> Given this realisation of the data (and conditioning on the true values of the reduced-form parameters), the conditional identified set for  $\varepsilon_{1t}$  is  $[-2.4, 2.6]$ , so the data are uninformative about the sign of the shock and its magnitude. The uniform distribution for  $\theta$  implies a distribution for  $\varepsilon_{1t}$  that assigns more mass to values towards the upper end of the identified set, and the 75th percentile of this distribution is around 2.4 standard deviations. The bound elicitation procedure in LMN would therefore suggest that it is reasonable to impose the shock-magnitude restriction that  $\varepsilon_{1t} \geq 2.4$ .



Notes:  $TN(\mu, 1)$  refers to truncated normal distribution for  $\theta$  with mean  $\mu$  and unit standard deviation. Dashed lines represent 75th percentiles of shock distribution induced by different distributions for  $\theta$ .

What if we used a different algorithm to randomly draw model parameters, corresponding to a different distribution for  $\theta$ ? To examine this question, I consider drawing  $\theta$  from a truncated normal distribution with mean parameter  $\mu$  and standard deviation one, where the lower and upper bounds are  $-\pi/2$  and  $\pi/2$ , respectively (i.e. the distribution is truncated to  $IS_\theta(\boldsymbol{\phi})$ ). I consider  $\mu \in \{-\pi/2, \pi/2\}$ , so the two distributions concentrate towards either end of  $IS_\theta(\boldsymbol{\phi})$ . When  $\mu = -\pi/2$ , the distribution for  $\varepsilon_{1t}$  assigns even more mass towards the upper end of the identified set, pushing up the 75th percentile relative to the case where  $\theta$  is uniformly distributed. When  $\mu = \pi/2$ , the distribution for  $\varepsilon_{1t}$  assigns more mass towards the lower end of the identified set, so the 75th percentile is slightly negative.

<sup>38</sup> This exercise assumes that  $\mathbf{y}_t = (1, -3)'$ . Given the value of  $\mathbf{H}$  in the DGP, this implies that  $\boldsymbol{\varepsilon}_t = (1.43, -2.14)'$ .

## Appendix B: Monte Carlo Exercises

This appendix details how the Monte Carlo exercises in Section 4.2 are implemented and describes the additional exercises that are referred to in the main text.

### B.1 Implementation details

In each Monte Carlo replication, the identified set for  $\theta$  under the baseline sign restrictions is computed analytically.<sup>39</sup> Under the shock-percentile restriction, the identified set is approximated using draws of  $\mathbf{Q}$  from the uniform distribution over  $\mathcal{Q}(\boldsymbol{\phi}|S)$ , obtained via the RWZ algorithm. At each draw of  $\mathbf{Q}$ , I compute the implied value of  $\theta$ , then approximate the identified set by the minimum and maximum values of  $\theta$  over the draws.<sup>40</sup> Montiel Olea and Nesbit (2021) provide results about the number of draws required to guarantee a given degree of approximation error. Their Theorem 3 provides an upper bound on the number of draws  $L$  from within the identified set required to guarantee that a misclassification error less than  $\varepsilon$  occurs with probability at least  $1 - \delta$ . The upper bound is  $\min\{2d \ln(2d/\delta), \exp(1)(2d + \ln(1/\delta))\}/\varepsilon$ , where  $d$  is the dimension of the parameter region being approximated. Setting  $d = 1$  (since  $\theta$  is scalar) and  $\varepsilon = \delta = 0.01$  yields an upper bound for  $L$  of around 1,060. I therefore approximate  $IS_{\theta}(\boldsymbol{\phi})$  using 1,060 draws of  $\mathbf{Q}$  in each Monte Carlo replication.

Because numerical approximation introduces error, the identified set obtained under the shock-percentile restriction can appear strictly smaller than under the baseline sign restrictions even when the shock-percentile restriction is not actually binding. When computing the share of Monte Carlo replications in which the shock-percentile restriction contributes no additional identifying information, I account for this approximation error by reporting the share of replications where the two identified sets differ in length by no more than 1 per cent.

### B.2 Additional exercises

The Monte Carlo exercise in Section 4.2 assumes a DGP where  $\text{vec}(\mathbf{H}) = (1, -0.3, 0.2, 1.2)'$ , which implies that  $p_t$  and  $q_t$  are weakly negatively correlated. Under this DGP, the shock-percentile restrictions often (sometimes substantially) sharpen identification. If the correlation were stronger, the identified set would be shorter under the baseline sign restrictions and the shock-percentile restrictions would tend to sharpen identification by less.

To illustrate this, I repeat the Monte Carlo exercise with an alternative DGP that yields tighter identification under the baseline. More specifically, the DGP now assumes that  $\text{vec}(\mathbf{H}) = (6, -1.8, 0.2, 1.2)'$ . Relative to the original DGP, these parameter values imply that the supply shock plays a much more important role in driving variation in  $p_t$  and  $q_t$ , inducing a stronger negative correlation between the two variables. Consequently, the identified set under the baseline sign restrictions is about 60 per cent shorter than the identified set under the original DGP.<sup>41</sup>

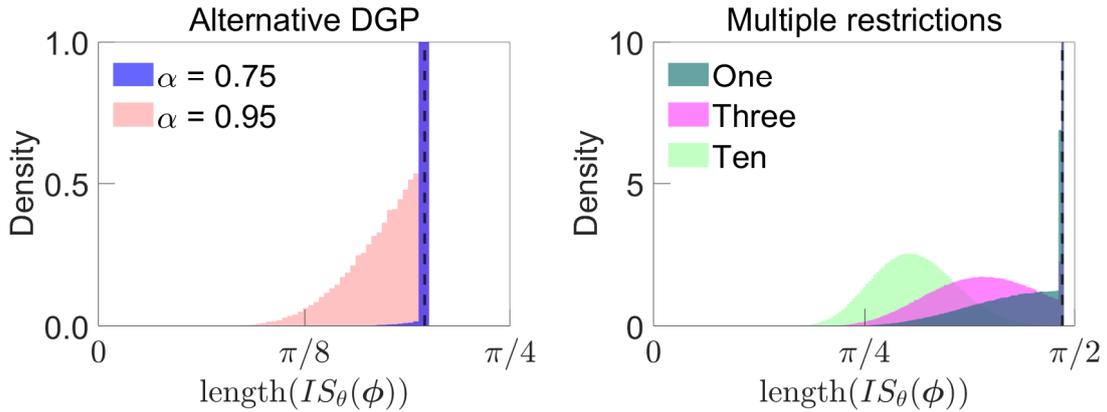
<sup>39</sup> For a characterisation of the identified set under these restrictions, see Baumeister and Hamilton (2015) or Read (2024).

<sup>40</sup> Draws of  $\mathbf{Q}$  are transformed into draws of  $\theta$  via  $\theta = \arccos(q_{11})\text{sign}(q_{21})$ , where  $q_{ij} = \mathbf{e}'_{i,2} \mathbf{Q} \mathbf{e}_{j,2}$ .

<sup>41</sup> It is well-known that the strength of the correlation between price and quantity innovations influences the sharpness of identification in this model of supply and demand. For example, see Uhlig (2017) or Read (2024).

The left panel of Figure B1 plots the Monte Carlo distribution of the length of the identified set under this alternative DGP. Under the weaker shock-percentile restriction ( $\alpha = 0.75$ ), the identified set is only about 0.1 per cent shorter, on average, than under the baseline sign restrictions, and the shock-percentile restriction is binding in only about 0.1 per cent of Monte Carlo replications. Even the stronger shock-percentile restriction ( $\alpha = 0.95$ ) only reduces the length of the identified set by about 1 per cent on average and does not sharpen identification at all in around 94 per cent of replications. Overall, the shock-percentile restrictions tend to be less informative than under the original DGP; they bind less frequently and, conditional on binding, tend to sharpen identification by less.

**Figure B1: Monte Carlo Distribution of Length of Identified Set under Alternative Assumptions**



Notes: Monte Carlo distribution of length of identified set  $IS_\theta(\phi)$  under combination of baseline and shock-percentile restrictions, where each Monte Carlo replication represents different realisation from DGP. Left panel assumes alternative DGP with sharper identification under baseline sign restrictions; right panel assumes original DGP and imposes multiple shock-percentile restrictions with  $\alpha = 0.75$ . Dashed line represents length of identified set under sign restrictions only. y-axes truncated for readability.

The Monte Carlo exercises in Section 4.2 and above impose only a single shock-percentile restriction. Imposing multiple restrictions – as in the empirical application in Section 5 – should tend to sharpen identification further. To illustrate this, I return to the original DGP considered in Section 4.2 and consider imposing multiple ‘weak’ shock-percentile restrictions with  $\alpha = 0.75$ . I consider imposing one, three or ten shock-percentile restrictions. In each exercise, the restrictions are imposed in the periods corresponding to the largest positive realisations of  $\varepsilon_{1t}$ ; that is, the shock-percentile restrictions take the form

$$\varepsilon_{1\tau}(\phi, \mathbf{Q}, \mathbf{u}_\tau) \geq G_1(0.75; \phi, \mathbf{Q}, \mathbf{U}) \quad \text{for } \tau \in \{t : \varepsilon_{1t} \geq \varepsilon_{1t}^{(r)}\} \quad (\text{B1})$$

where  $\varepsilon_{1t}^{(r)}$  is the  $r$ th largest order statistic of the shock sequence  $\{\varepsilon_{1t}\}_{t=1}^T$ . In the cases considered,  $r \in \{1, 3, 10\}$ .

The right panel of Figure B1 plots the Monte Carlo distribution of the length of the identified set for each set of restrictions. It is apparent that adding additional shock-percentile restrictions can substantially sharpen identification; the Monte Carlo distribution of the length of the identified set shifts to the left as more shock-percentile restrictions are added. With three shock-percentile restrictions, the length of the identified set is 18 per cent shorter, on average, than under the baseline sign restrictions (compared with 7 per cent when there is one restriction). Under ten restrictions the length of the identified set is 36 per cent shorter, on average, than under the baseline sign restrictions.

## Appendix C: Additional Details about Empirical Exercises

This appendix provides additional details related to the empirical applications. Appendix C.1 describes the numerical procedures used to approximate identified sets in both applications. I then revisit two exercises from LMN using the shock-percentile restrictions: Appendix C.2 examines whether recursive identification schemes are consistent with the data; and Appendix C.3 examines whether the structural shocks are non-Gaussian.

### C.1 Numerical implementation

To approximate identified sets in Section 5, I obtain draws of  $\mathbf{Q}$  from (an approximation of) the uniform distribution over its identified set, compute the parameters of interest, and take the minimum and maximum of the parameter over the draws of  $\mathbf{Q}$ .<sup>42</sup> Given a finite number of draws, this strategy will generate identified sets that are too narrow, with the approximation error vanishing as the number of draws increases. As discussed in Section B.1, Montiel Olea and Nesbit (2021) provide results about the number of draws required to guarantee a given degree of approximation error. Based on the upper bound in their Theorem 3, approximating an identified set of dimension  $d = 3 \times 3 \times 61 = 549$  with misclassification error less than  $\varepsilon = 0.01$  and probability at least  $1 - \delta = 0.99$  requires  $L \approx 300,000$  draws from inside the identified set.

Under the event and external variable constraints considered in LMN, the identified set for  $\mathbf{Q}$  is substantially truncated relative to the unconstrained identified set. Obtaining a sufficient number of draws satisfying the restrictions via accept-reject sampling is consequently extremely computationally burdensome. I therefore use the sampler based on ‘soft sign restrictions’ proposed in Read and Zhu (2025), which can be more computationally efficient than accept-reject sampling in cases like this where identification is ‘tight’.

The sampler uses Markov chain Monte Carlo methods (specifically, the slice sampler) to obtain draws of  $\mathbf{Q}$  from a target distribution that is similar to the uniform distribution but that smoothly penalises parameter values that violate the identifying restrictions. This yields some draws that violate the restrictions, and draws that satisfy the restrictions are not uniformly distributed. An importance-sampling step can be used to discard the draws that violate the restrictions and resample the remaining draws so that they better approximate the target uniform distribution. However, since the objective here is only to approximate the bounds of the identified set, it suffices to discard the draws that violate the identifying restrictions, leaving only draws from inside the identified set, and the resampling step is unnecessary. Otherwise, the implementation of the sampler follows Read and Zhu (2025).<sup>43</sup>

To obtain  $L$  effective draws from inside the identified set, it is necessary to obtain more than  $L$  draws from the smoothed target density, because some draws will not satisfy the identifying restrictions. For each set of identifying restrictions, I therefore gross up  $L$  by an initial estimate of the effective sample size for the sampler. For example, under the restrictions from LMN, the effective sample size is around 89 per cent, so I obtain  $L/0.89$  draws, which yields approximately  $L$  draws from inside the identified set.

<sup>42</sup> This approach follows Algorithm 2 in Giacomini and Kitagawa (2021a).

<sup>43</sup> The sampler from Read and Zhu (2025) requires specifying a tuning parameter  $\Delta$  that determines how strongly parameter values are penalised when they violate the identifying restrictions. I set  $\Delta = 10^{-5}$ .

To implement the robust Bayesian approach to inference in the empirical application in Section 6, I use the same approach to approximating identified sets at each draw of the reduced-form parameters. Based on the upper bound from Theorem 3 of Montiel Olea and Nesbit (2021), approximating an identified set of dimension  $d = 2 \times 61 = 122$  with misclassification error less than  $\varepsilon = 0.01$  and probability at least  $1 - \delta = 0.99$  requires  $L \approx 13,000$  draws of  $\mathbf{Q}$  at each draw of the reduced-form parameters.<sup>44</sup> The results are based on 1,000 draws of the reduced-form parameters.

## C.2 Recursive identification

Much of the empirical literature on the macroeconomic effects of uncertainty relies on recursive identifying schemes to estimate the effects of uncertainty shocks.<sup>45</sup> LMN demonstrate that recursive structures are inconsistent with their results. This section examines whether this remains the case when replacing the shock-magnitude restrictions with shock-percentile restrictions.

There are three variables in the SVAR, so there are  $3! = 6$  possible recursive orderings. Let

$$\mathbf{H} = \begin{bmatrix} H_{MM} & H_{MY} & H_{MF} \\ H_{YM} & H_{YY} & H_{YF} \\ H_{FM} & H_{FY} & H_{FF} \end{bmatrix} \quad (\text{C1})$$

so  $H_{ij}$  is the impact response of variable  $i$  to shock  $j$ . The six possible recursive orderings are:

$$\begin{aligned} & \begin{bmatrix} H_{MM} & 0 & 0 \\ H_{YM} & H_{YY} & 0 \\ H_{FM} & H_{FY} & H_{FF} \end{bmatrix}, \quad \begin{bmatrix} H_{MM} & 0 & 0 \\ H_{FM} & H_{FF} & 0 \\ H_{YM} & H_{YF} & H_{YY} \end{bmatrix}, \quad \begin{bmatrix} H_{YY} & 0 & 0 \\ H_{MY} & H_{MM} & 0 \\ H_{FY} & H_{FM} & H_{FF} \end{bmatrix}, \\ & \begin{bmatrix} H_{YY} & 0 & 0 \\ H_{FY} & H_{FF} & 0 \\ H_{MY} & H_{MF} & H_{MM} \end{bmatrix}, \quad \begin{bmatrix} H_{FF} & 0 & 0 \\ H_{MF} & H_{MM} & 0 \\ H_{YF} & H_{YM} & H_{YY} \end{bmatrix}, \quad \begin{bmatrix} H_{FF} & 0 & 0 \\ H_{YF} & H_{YY} & 0 \\ H_{MF} & H_{MY} & H_{MM} \end{bmatrix} \end{aligned} \quad (\text{C2})$$

At a given value of  $\phi$ , we can assess whether a specific recursive ordering is consistent with the data and identifying restrictions by: 1) computing the (point-identified) structural parameters under the recursive ordering; and 2) checking whether the event and external variable constraints are satisfied (i.e. that the identified set is non-empty).<sup>46</sup> If none of the structural parameters implied by the six recursive orderings are consistent with the event and external variable constraints, recursive structures are inconsistent with the data (given the other identifying restrictions).

At the OLS estimates of  $\phi$ , none of the possible recursive orderings are consistent with the event and external variable constraints used in LMN. This continues to be the case when replacing the shock-magnitude restrictions with the shock-percentile restrictions (with  $\alpha = 0.75$ ). This suggests that recursive orderings are not supported by the data, conditional on the other identifying restrictions. The results under the shock-percentile restrictions therefore reinforce the challenges associated with

<sup>44</sup> I present identified sets for the responses of two variables (the federal funds rate and output) for horizons  $h = 0, 1, \dots, 60$ . Approximating identified sets for the responses of additional variables with the same degree of accuracy would require more draws of  $\mathbf{Q}$ .

<sup>45</sup> See Kilian, Plante and Richter (2025) for many examples.

<sup>46</sup> LMN check whether the identified sets for the relevant elements of  $\mathbf{H}$  exclude zero, which is a sufficient condition for the identified sets to be inconsistent with a recursive ordering.

relying on recursive assumptions when attempting to tease apart the causal relationship between uncertainty and real activity.<sup>47</sup>

### C.3 Are the shocks non-Gaussian?

Focusing on a particular set of parameters within the identified set, LMN document that the implied structural shocks exhibit features of non-Gaussianity, such as skewness and excess kurtosis. This result may be of interest because researchers often assume that structural shocks are normally distributed when using Bayesian methods to estimate SVARs, and departures from Gaussianity can distort inference (e.g. Petrova 2022). This section examines whether evidence of non-Gaussianity remains under the shock-percentile restrictions.

Rather than focusing on the properties of a single sequence of shocks implied by a particular set of parameters, as in LMN, I compute identified sets for the skewness and kurtosis of each structural shock. Intuitively, each  $\mathbf{Q}$  in the identified set is associated with a different sequence of shocks, and thus (potentially) a different skewness and kurtosis. I characterise identified sets for the skewness and kurtosis of the shocks by randomly drawing values of  $\mathbf{Q}$  from the uniform distribution over  $\mathcal{Q}(\phi|S)$ , computing the skewness and kurtosis of the implied shocks, and taking the minimum and maximum of these statistics over the draws of  $\mathbf{Q}$ .<sup>48</sup>

Table C1 presents identified sets for the skewness and kurtosis of the structural shocks under the identifying restrictions from LMN and the shock-percentile restrictions with  $\alpha = 0.75$ . If the shocks were Gaussian, we should expect them to have skewness equal to zero and kurtosis equal to three. The identified sets under the LMN restrictions unambiguously point to the presence of non-Gaussian features in the identified shocks; macroeconomic uncertainty shocks are positively skewed, output and financial uncertainty shocks are negatively skewed, and all shocks display excess kurtosis.

<b>Table C1: Identified Sets for Skewness and Kurtosis of Structural Shocks</b>				
<b>Shock</b>	<b>Skewness</b>		<b>Kurtosis</b>	
	LMN	SP <sup>(a)</sup>	LMN	SP <sup>(a)</sup>
$\varepsilon_{Mt}$	[0.30, 0.56]	[-0.70, 0.56]	[5.56, 5.90]	[5.34, 10.63]
$\varepsilon_{Yt}$	[-0.47, -0.28]	[-1.01, -0.03]	[5.38, 5.97]	[4.84, 11.03]
$\varepsilon_{Ft}$	[-1.69, -1.48]	[-1.88, 0.11]	[16.87, 19.16]	[7.33, 20.95]

Note: (a) 'SP' corresponds to shock-percentile restrictions with  $\alpha = 0.75$ .

While the identified sets under the shock-percentile restrictions are wider than under the LMN restrictions, they continue to suggest that the structural shocks possess some non-Gaussian features. The identified sets for skewness contain zero for the uncertainty shocks and values close to zero for the output shock, so there is no longer clear evidence of skewness in the identified shocks. However, the identified sets unambiguously point to excess kurtosis in all three shocks.

47 Kilian *et al* (2025) demonstrate that this challenge can not be overcome by examining the results obtained under alternative causal orderings, because these do not necessarily bound the true impulse responses when the DGP does not feature a recursive structure.

48 A similar idea underlies the approach to identification in Andrade *et al* (2024). Given that higher-order moments are set identified under sign restrictions, they impose inequality constraints on higher-order moments to sharpen identification.

## References

- Abbate A, S Eickmeier and E Prieto (2023)**, 'Financial Shocks and Inflation Dynamics', *Macroeconomic Dynamics*, 27(2), pp 350–378.
- Andrade P, F Ferroni and L Melosi (2024)**, 'Higher-order Moment Inequality Restrictions for SVARs', Centre for Economic Policy Research Discussion Paper DP19813.
- Andreasen MM, G Caggiano, E Castelnuovo and G Pellegrino (2024)**, 'Does Risk Matter More in Recessions than in Expansions? Implications for Monetary Policy', *Journal of Monetary Economics*, 143, Article 103533.
- Angelini G, E Bacchiocchi, G Caggiano and L Fanelli (2019)**, 'Uncertainty across Volatility Regimes', *Journal of Applied Econometrics*, 34(3), pp 437–455.
- Antolín-Díaz J and JF Rubio-Ramírez (2018)**, 'Narrative Sign Restrictions for SVARs', *The American Economic Review*, 108(10), pp 2802–2829.
- Arias JE, JF Rubio-Ramírez and DF Waggoner (2018)**, 'Inference Based on Structural Vector Autoregressions Identified with Sign and Zero Restrictions: Theory and Applications', *Econometrica*, 86(2), pp 685–720.
- Arias JE, JF Rubio-Ramírez and DF Waggoner (2025)**, 'Uniform Priors for Impulse Responses', *Econometrica*, 93(2), pp 695–718.
- Baker SR, N Bloom and SJ Terry (2024)**, 'Using Disasters to Estimate the Impact of Uncertainty', *The Review of Economic Studies*, 91(2), pp 720–747.
- Baumeister C and JD Hamilton (2015)**, 'Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information', *Econometrica*, 83(5), pp 1963–1999.
- Baumeister C and JD Hamilton (2018)**, 'Inference in Structural Vector Autoregressions when the Identifying Assumptions Are Not Fully Believed: Re-evaluating the Role of Monetary Policy in Economic Fluctuations', *Journal of Monetary Economics*, 100, pp 48–65.
- Baumeister C and JD Hamilton (2019)**, 'Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks', *The American Economic Review*, 109(5), pp 1873–1910.
- Baumeister C and JD Hamilton (2024)**, 'Advances in Using Vector Autoregressions to Estimate Structural Magnitudes', *Econometric Theory*, 40(3), pp 472–510.
- Ben Zeev N (2018)**, 'What Can We Learn about News Shocks from the Late 1990s and Early 2000s Boom-bust Period?', *Journal of Economic Dynamics & Control*, 87, pp 94–105.
- Bloom N (2009)**, 'The Impact of Uncertainty Shocks', *Econometrica*, 77(3), pp 623–685.
- Braun R (2023)**, 'The Importance of Supply and Demand for Oil Prices: Evidence from Non-Gaussianity', *Quantitative Economics*, 14(4), pp 1163–1198.
- Brianti M (2025)**, 'Financial Shocks, Uncertainty Shocks, and Corporate Liquidity', *Journal of Applied Econometrics*, 40(7), pp 814–828.

- Caggiano G and E Castelnuovo (2023)**, 'Global Financial Uncertainty', *Journal of Applied Econometrics*, 38(3), pp 432–449.
- Caggiano G, E Castelnuovo and N Groshenny (2014)**, 'Uncertainty Shocks and Unemployment Dynamics in U.S. Recessions', *Journal of Monetary Economics*, 67, pp 78–92.
- Caggiano G, E Castelnuovo and R Kima (2020)**, 'The Global Effects of Covid-19-induced Uncertainty', *Economics Letters*, 194, Article 109392.
- Caggiano G, E Castelnuovo, R Kima and S Delrio (2020)**, 'Financial Uncertainty and Real Activity: The Good, the Bad, and the Ugly', Australian National University Crawford School of Public Policy, Centre for Applied Macroeconomic Analysis, CAMA Working Paper 67/2020.
- Carriero A, M Marcellino and T Tornese (2023)**, 'Macro Uncertainty in the Long Run', *Economics Letters*, 225, Article 111067.
- Ferreira LN (2022)**, 'Forward Guidance Matters: Disentangling Monetary Policy Shocks', *Journal of Macroeconomics*, 73, Article 103423.
- Fry R and A Pagan (2011)**, 'Sign Restrictions in Structural Vector Autoregressions: A Critical Review', *Journal of Economic Literature*, 49(4), pp 938–960.
- Giacomini R and T Kitagawa (2021a)**, 'Robust Bayesian Inference for Set-identified Models', *Econometrica*, 89(4), pp 1519–1556.
- Giacomini R and T Kitagawa (2021b)**, 'Supplement to "Robust Bayesian Inference for Set-identified Models"', *Econometrica*, 89(4), Supporting Information, Online Appendix.
- Giacomini R, T Kitagawa and M Read (2021)**, 'Identification and Inference under Narrative Restrictions', Unpublished manuscript, February. Available at <<https://doi.org/10.48550/arXiv.2102.06456>>.
- Giacomini R, T Kitagawa and M Read (2023)**, '[Identification and Inference under Narrative Restrictions](#)', RBA Research Discussion Paper No 2023-07.
- Himounet N (2022)**, 'Searching the Nature of Uncertainty: Macroeconomic and Financial Risks VS Geopolitical and Pandemic Risks', *International Economics*, 170, pp 1–31.
- Inoue A and L Kilian (2025)**, 'The Conventional Impulse Response Prior in VAR Models with Sign Restrictions', Centre for Economic Policy Research Discussion Paper DP20159.
- Inoue A and L Kilian (2026)**, 'When Is the Use of Gaussian-inverse Wishart-Haar Priors Appropriate?', *Journal of Political Economy*, 134(2), pp 773–794.
- Jurado K, SC Ludvigson and S Ng (2015)**, 'Measuring Uncertainty', *The American Economic Review*, 105(3), pp 1177–1216.
- Kang S and K Park (2024)**, 'Endogenous Uncertainty and Monetary Policy', *Macroeconomic Dynamics*, 28(2), pp 462–477.
- Kilian L and H Lütkepohl (2017)**, *Structural Vector Autoregressive Analysis*, Themes in Modern Econometrics, Cambridge University Press, Cambridge.

- Kilian L, MD Plante and AW Richter (2025)**, 'Macroeconomic Responses to Uncertainty Shocks: The Perils of Recursive Orderings', *Journal of Applied Econometrics*, 40(4), pp 395–410.
- Lanne M, M Meitz and P Saikkonen (2017)**, 'Identification and Estimation of Non-Gaussian Structural Vector Autoregressions', *Journal of Econometrics*, 196(2), pp 288–304.
- Leduc S and Z Liu (2016)**, 'Uncertainty Shocks Are Aggregate Demand Shocks', *Journal of Monetary Economics*, 82, pp 20–35.
- Lowenstein R (2000)**, *When Genius Failed: The Rise and Fall of Long-Term Capital Management*, Random House, New York.
- Ludvigson SC, S Ma and S Ng (2017)**, 'Shock Restricted Structural Vector-autoregressions', NBER Working Paper No 23225, rev January 2020.
- Ludvigson SC, S Ma and S Ng (2021)**, 'Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?', *American Economic Journal: Macroeconomics*, 13(4), pp 369–410.
- Manski CF (2003)**, *Partial Identification of Probability Distributions*, Springer Series in Statistics, Springer-Verlag, New York.
- Manski CF (2011)**, 'Policy Analysis with Incredible Certitude', *The Economic Journal*, 121(554), pp F261–F289.
- Mertens K and MO Ravn (2013)**, 'The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States', *The American Economic Review*, 103(4), pp 1212–1247.
- Montiel Olea JL and J Nesbit (2021)**, '(Machine) Learning Parameter Regions', *Journal of Econometrics*, 221(1, Part C), pp 716–744.
- Moon HR and F Schorfheide (2012)**, 'Bayesian and Frequentist Inference in Partially Identified Models', *Econometrica*, 80(2), pp 755–782.
- Petrova K (2022)**, 'Asymptotically Valid Bayesian Inference in the Presence of Distributional Misspecification in VAR Models', *Journal of Econometrics*, 230(1), pp 154–182.
- Poirier DJ (1998)**, 'Revising Beliefs in Nonidentified Models', *Econometric Theory*, 14(4), pp 483–509.
- Ramey VA (2016)**, 'Macroeconomic Shocks and Their Propagation', in JB Taylor and H Uhlig (eds), *Handbook of Macroeconomics: Volume 2A*, Handbooks in Economics, Elsevier, Amsterdam, pp 71–162.
- Read M (2022)**, 'Algorithms for Inference in SVARs Identified With Sign and Zero Restrictions', *The Econometrics Journal*, 25(3), pp 699–718.
- Read M (2024)**, '[Sign Restrictions and Supply-demand Decompositions of Inflation](#)', RBA Research Discussion Paper No 2024-05.
- Read M (forthcoming)**, 'Set-identified Structural Vector Autoregressions and the Effects of a 100 Basis Point Monetary Policy Shock', *The Review of Economics and Statistics*.

**Read M and D Zhu (2025)**, 'Fast Posterior Sampling in Tightly Identified SVARs Using "Soft" Sign Restrictions', RBA Research Discussion Paper No 2025-03.

**Romer CD and DH Romer (1989)**, 'Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz', in OJ Blanchard and S Fischer (eds), *NBER Macroeconomics Annual 1989*, 4, The MIT Press, Cambridge, pp 121–170.

**Rubio-Ramírez JF, DF Waggoner and T Zha (2010)**, 'Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference', *The Review of Economic Studies*, 77(2), pp 665–696.

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

**Theophilopoulou A (2022)**, 'The Impact of Macroeconomic Uncertainty on Inequality: An Empirical Study for the United Kingdom', *Journal of Money, Credit and Banking*, 54(4), pp 859–884.

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

**Uhlig H (2017)**, 'Shocks, Sign Restrictions, and Identification', in B Honoré, A Pakes, M Piazzesi and L Samuelson (eds), *Advances in Economics and Econometrics: Eleventh World Congress: Volume 2*, Econometric Society Monographs, Cambridge University Press, Cambridge, pp 95–127.

**Watson MW (2020)**, 'Comment on "On the Empirical (Ir)relevance of the Zero Lower Bound"', in MS Eichenbaum, E Hurst and JA Parker (eds), *NBER Macroeconomics Annual 2019*, 34, University of Chicago Press, Chicago, pp 182–193.