Refining Set-Identification in VARs through Independence

Thorsten Drautzburg and Jonathan H. Wright*

August 2021

Abstract

Identification in VARs has traditionally mainly relied on second moments. Some researchers have considered using higher moments as well, but there are concerns about the strength of the identification obtained in this way. In this paper, we propose refining existing identification schemes by augmenting sign restrictions with a requirement that rules out shocks whose higher moments significantly depart from independence. This approach does not assume that higher moments help with identification; it is robust to weak identification. In simulations we show that it controls coverage well, in contrast to approaches that assume that the higher moments deliver point identification. However, it requires large sample sizes and/or considerable non-normality to reduce the width of confidence intervals by much. We consider some empirical applications. We find that it can reject many possible rotations. The resulting confidence sets for impulse responses may be non-convex, corresponding to disjoint parts of the space of rotation matrices. We show that in this case, augmenting sign and magnitude restrictions with an independence requirement can yield bigger gains.

Keywords: vector-autoregression, sign restrictions, set identification, weak identification, independent shocks, non-convex confidence set.

JEL codes: C32; C51.

1 Introduction

Vector Autoregressive (VAR) models have been one of the workhorse models in macroeconomics since Sims (1980). They make efficient use of second moments for semi-structural analysis in macroeconomics. Traditionally, the focus has been on point-identification of one or more shocks in the VAR, via short-run restrictions (Sims, 1980), long-run restrictions (Blanchard and Quah, 1989), or the max-share approach at intermediate frequencies (Uhlig, 2003). More recently, there has been an interest in exploring sign restrictions on responses to shocks (Uhlig, 1998; Faust, 1998; Canova and De Nicolo, 2002; Uhlig, 2005) for identification. Sign restrictions no longer point-identify shocks, often yielding large identified sets even when imposing prior information over the non-identifiable parameters (Kilian and Murphy, 2012).

This paper proposes an approach to test point-identified VARs and refine the frequentist identified set with sign restrictions. The VAR identification problem is one of deconvolution. The researcher observes p forecast errors $\mathbf{u}_t$ and their covariance $\boldsymbol{\Sigma}$. These forecast errors are related to the shocks of interest $\boldsymbol{\varepsilon}_t$ by some linear map $\mathbf{u}_t = \mathbf{A}\boldsymbol{\varepsilon}_t$. Shocks are assumed to be independently and continuously distributed with unit variances: $\boldsymbol{\varepsilon}_t \sim ID(\mathbf{0}, \mathbf{I}_q)$. If researchers use only the first two moments or if shocks are normally distributed, deconvoluting the forecast errors $\mathbf{u}_t$ to yield the shocks of interest is not possible: Any matrix $\tilde{\mathbf{A}} \equiv \text{chol}(\Sigma)\mathbf{Q}$, where $\mathbf{Q} \in \mathcal{O}^q$ is an orthonormal matrix, is observably equivalent. However, the insight of Independent Component Analysis (ICA, see, for example, Hyvärinen et al. (2004)) is that higher moments allow one to solve such deconvolution problems when the underlying shocks are non-Gaussian. This paper proposes a more robust approach: It simply proposes estimating the confidence region for impulse-response functions (IRFs) as the intersection of the point-identified or set-identified confidence interval based on sign restrictions and a region that rules out shocks whose higher moments exhibit significant co-dependence. Naturally, this method requires that the true structural shocks are independent, not just uncorrelated. This however seems reasonable if we think of the structural shocks as primitive underlying causes of fluctuations (Bernanke, 1986).

The proposed robust approach can yield point-identification in large samples, but remains well defined in the Gaussian limit. Specifically, the estimator for impulse-response function(s) of interest is the identified set based on sign restrictions intersected with an inverted Generalized Method of Moments (GMM) test statistic based on co-dependence of higher moments or a non-parametric test of independence (Hoeffding, 1948; Blum et al., 1961). In population, these criteria rule out incorrect identification schemes when the data are non-Gaussian. One of our GMM test statistics allows for stochastic volatility while disallowing for other higher-order dependence.

When the data is Gaussian, there is no co-dependence of higher moments and the estimator becomes close to the plug-in estimator in Moon and Schorfheide (2012): Any $\mathbf{Q}$ results in IRFs consistent with the restrictions. When the sample size is small, the inverted test statistic rules out little of $\mathcal{O}^q$ and the refinement again has little impact. With significant non-Gaussianity and sufficiently large samples, however, one can rule out large subsets of $\mathcal{O}^q$ and materially refine the

identified set – or reject the proposed identification scheme in the case of point-identification.[1] Our paper exploits independence for identification, but does so allowing for weak identification, as called for by Montiel Olea et al. (2021).

The identification that we obtain from the independence requirement is inherently local, not global. This is because if any matrix $\mathbf{Q}$ satisfies the independence requirement, then so does another matrix that consists of permuting the columns of $\mathbf{Q}$ and/or flipping their signs. This is why we consider imposing the independence requirement as a refinement of sign restrictions, not as the sole source of identification.

We use a Bonferroni approach for inference, building on Granziera et al. (2018). While conservative, this is more efficient and computationally feasible than a projection approach and it prevents sampling uncertainty in the estimation of higher moments from contaminating the estimates of second moments. We thus invert our test statistics for dependence separately, using the bootstrap to account for sampling uncertainty induced by the estimation of reduced form parameters. Due to the local nature of statistical identification, this estimator can yield non-convex confidence intervals.

This paper has three parts: First, it shows that for a class of data-generating processes, higher moments exhibit zero co-dependence. This observation motivates a bootstrap-based test statistic, whose inversion yields a confidence interval for $\mathbf{Q}$. Its intersection with the identified set for $\mathbf{Q}$ in terms of the sign restrictions is the proposed estimator, and is a refinement of the standard sign restrictions identified set. We also consider another bootstrap-based test statistic based on the test for independence of Hoeffding (1948) and Blum et al. (1961). Second, we illustrate how the proposed approach works using simulations with various Gaussian and non-Gaussian shocks. Third, we consider empirical applications. We consider testing of existing point identification schemes: we revisit the VARs of Blanchard and Quah (1989) and Bloom (2009), as well as the Smets and Wouters (2007) DSGE model, and find that independence of the shocks is rejected in many cases. We also consider how our method applies to inference for confidence intervals in a structural VAR in employment and wages. We find that our refinement eliminates many possible rotations, but for the labor demand shock it makes the confidence set non-convex because there are pairs of orthogonal rotations that both satisfy the sign and independence conditions, in another manifestation of the local nature of identification. We show that if the independence refinement is applied in conjunction with both sign and magnitude restrictions, then it achieves a large reduction in the size of the identified set.

Literature. Lanne et al. (2017), Lanne and Luoto (2019), Gouriéroux et al. (2017) and Guay (2021) also use insights from ICA to identify VARs. In their book on VARs, Kilian and Lütkepohl (2018) treat the use of non-normality as an important potential strategy for identification. All of this work, however, focuses on point identification via higher moments and requires non-Gaussian shocks. In contrast to this approach of achieving statistical identification, we rely on achieving identification by economic assumptions alone and then to use higher moments test or refine

---

[1]In theory, our approach could also reject the identifying assumptions for set identification. But with empirically relevant sample sizes and data generating processes, this possibility is not relevant.

the identification based on those economic assumptions. The key advantage of our approach is that it allows for Gaussian shocks and is thus robust to lack of identification. Lanne and Luoto (2020) show that, under t-distributed shocks, higher moments identify shocks and allow for sign restrictions to be tested. This note is complementary: It takes the sign restrictions as given, but proposes an estimator that is informed by possible departures from Gaussianity.

Several papers propose frequentist estimators for set-identified impulse-response functions, using the Delta method (Gafarov et al., 2018), using a projection method (Gafarov et al., 2016), and using a Bonferroni approach (Granziera et al., 2018). Here, we follow the Bonferroni approach of Granziera et al. (2018) because it prevents inaccuracies in the estimation of higher moments from affecting the estimation of second moments. This improves our estimator when higher moments are only imprecisely estimated. In addition, the estimator can recover non-convex confidence intervals. Bacchiocchi and Kitagawa (2020) also propose an estimator capable of recovering non-convex confidence intervals, but they do not consider statistical identification.

Kilian and Murphy (2012) and Amir-Ahmadi and Drautzburg (2021) show how restrictions on linear combinations of sign restrictions can narrow the identified set. This paper shares their goal, but provides a complementary approach: Given a set of inequality restrictions, we show that restrictions on co-skewness and co-kurtosis can narrow the identified set.

A related idea to identification from non-normality is identification through time-varying volatility, considered by Rigobon (2003) and Lewis (2021). This is a similar approach in that it goes beyond unconditional second moments to help with identification, but is not the focus of this paper.

The plan for the remainder of this paper is as follows. In section 2, we define the structural VAR and the proposed approach to identification. Sections 3 and 4 report Monte Carlo simulations and empirical applications. Section 5 concludes.

2 Model

2.1 Model setup and definitions

Let $\mathbf{y}_t$ be a $p \times 1$ vector of observables. A VAR(k), with k lags, is defined as:

$$\mathbf{y}_t = \mathbf{b}_0 + \sum_{\ell=1}^{k} \mathbf{B}_\ell \mathbf{y}_{t-\ell} + \mathbf{u}_t, \qquad \mathbf{u}_t = \mathbf{A}\boldsymbol{\varepsilon}_t, \qquad \mathrm{Cov}[\mathbf{u}_t] = \mathbf{A}\mathbf{A}' \equiv \boldsymbol{\Sigma}. \tag{2.1}$$

Here, $\mathbf{b}_0$ characterizes the conditional mean of $\mathbf{y}_t$, while $\mathbf{B}_\ell$ captures lagged dynamics. $\mathbf{u}_t$ is the vector of forecast errors with full-rank covariance matrix $\boldsymbol{\Sigma} \equiv \mathbb{E}[\mathbf{u}_t\mathbf{u}_t'] = \mathbf{A}\mathbf{A}'$. $\mathbf{A}$ is the *impact matrix*. Let $\boldsymbol{\Sigma}^{\mathrm{tr}}$ be the Cholesky factor of $\boldsymbol{\Sigma}$. $\boldsymbol{\varepsilon}_t$ represents the p dimensional vector of *structural shocks*.

The identification problem consists of recovering $\boldsymbol{\varepsilon}_t$ from $\mathbf{u}_t$. Equivalently, it requires backing out $\mathbf{A}$ from $\boldsymbol{\Sigma}$. In general, any $\tilde{\mathbf{A}} = \boldsymbol{\Sigma}^{\mathrm{tr}}\mathbf{Q}$, $\mathbf{Q}\mathbf{Q}' = \mathbf{I}_p$ is observationally equivalent when considering

(first and) second moments only: $\tilde{\mathbf{A}}\tilde{\mathbf{A}}' = \boldsymbol{\Sigma}^{\text{tr}}\mathbf{Q}\mathbf{Q}'\boldsymbol{\Sigma}^{\text{tr}'} = \boldsymbol{\Sigma}$. Associated with each $\tilde{\mathbf{A}}$ and given the vector of forecast errors $\mathbf{u}_t$ is an estimate of the underlying structural shocks:

$$\hat{\varepsilon}_t(\tilde{\mathbf{A}}) \equiv \tilde{\mathbf{A}}^{-1}\mathbf{u}_t, \tag{2.2a}$$

$$= \tilde{\mathbf{A}}^{-1}\mathbf{A}\boldsymbol{\varepsilon}_t. \tag{2.2b}$$

This estimate of the structural shocks is a key ingredient of proposed procedure going forward.

Given any $\mathbf{Q}$, an *impulse-response function* (IRF) $\{\boldsymbol{\phi}_{i,s}(\mathbf{Q})\}_{s=0}^{\infty}$ is the sequence of forecast revisions given a unit shock $\epsilon_{it}, i \in \{1, \ldots, p\}$. In a VAR(1), the IRF is the sequence $\boldsymbol{\phi}_{i,s}\mathbf{B}_1^s\boldsymbol{\Sigma}^{\text{tr}}\mathbf{Q}_{\circ,i}$. This definition readily generalizes to a VAR with k lags using its companion form representation.

Let $\mathbf{W}_i, i = 1\ldots, p$ denote the r_i restrictions imposed on the ith shock. We say that $\mathbf{Q}$ satisfies the *sign restrictions* if $\mathbf{W}_i\mathbf{Q} \geq 0, i = 1, \ldots, p$. The *identified set for* $\mathbf{Q}$ is thus:

$$\mathcal{Q} = \{\mathbf{Q} \in \mathcal{O}^p | \mathbf{W}_i\mathbf{Q} \geq 0, i = 1, \ldots, p\}. \tag{2.3}$$

Given $\{\mathbf{B}_\ell\}, \boldsymbol{\Sigma}$, the associated *identified set for* $\boldsymbol{\phi}_{\circ,s}$, the IRFs, is the image $\mathcal{F}_s = f_s(\mathcal{Q})$, where $f(\mathbf{Q}) = [\boldsymbol{\phi}_{1,s}(\mathbf{Q}), \ldots, \boldsymbol{\phi}_{p,s}(\mathbf{Q})] : \mathcal{O}^p \to \mathbb{R}^{p \times p}$ is a linear function.

Conceptually, the analysis below can proceed in the static case of a VAR(0). In practice, however, the dynamics allow researchers to impose additional restrictions on the shocks via restrictions on the IRFs.

While it is common to assume that $\boldsymbol{\varepsilon}_t \overset{iid}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$, this note proceeds without parametric distributional assumptions.

Assumption 1 (Strong independence)**.** Each ϵ_i is independently and continuously distributed with mean zero, unit variance, and finite moments up to order $\bar{m}$:

$$\boldsymbol{\varepsilon} \sim F_1 \times \ldots \times F_p, \tag{2.4a}$$

$$\forall i = 1, \ldots, p: \quad \mathbb{E}[\epsilon_i] = 0, \quad \mathbb{E}[\epsilon_i^2] = 1, \tag{2.4b}$$

$$\forall i = 1, \ldots, p, m = 1, \ldots, \bar{m} : \mathbb{E}[\epsilon_i^{\bar{m}}] < \infty. \tag{2.4c}$$

One concern about requiring independence and identical distributions is that it rules out stochastic volatility. The assumption of identical distributions – and thus of constant volatility – may be relaxed in the context of the independence assumption as recently shown by Gaigall (forthcoming) in the bivariate case. Here, however, we use a moment-based approach to accommodate many forms of stochastic volatility. This results in the following, weaker condition in Assumption 2:

Assumption 2 (Conditional independence)**.** Each ϵ_i has mean zero, unit variance, and is uncor-

related with higher powers of other shocks. All moments up to order $\bar{m}$ are finite:

$$\forall i = 1, \ldots, p: \quad \mathbb{E}[\epsilon_i] = 0, \quad \mathbb{E}[\epsilon_i^2] = 1, \tag{2.5a}$$

$$\forall i = 1, \ldots, p, j \neq i, m = 1, \ldots, \bar{m}: \quad \mathbb{E}[\epsilon_i \epsilon_j^m] = 0, \tag{2.5b}$$

$$\forall i = 1, \ldots, p, m = 1, \ldots, \bar{m}: \mathbb{E}[\epsilon_i^{\bar{m}}] < \infty. \tag{2.5c}$$

For $m = 3$, (2.5b) is one of the assumptions in Lanne and Luoto (2019).

2.2 Detecting independence

We now discuss two approaches for detecting whether shocks are independent. First, we consider a criterion based on third and fourth moments. This criterion has intuitive appeal – but may fail to detect some forms of dependence. Second, we consider a non-parametric, direct test of the independence assumption (2.4a).

2.2.1 Dependence of higher order moments

To quantify dependence of higher order moments, we focus on co-skewness and co-kurtosis. These are third and fourth moment analogues of the correlation. The existing literature on using higher order moments for identification has used these (e.g.. Guay (2021) uses third and fourth moments, while Lanne and Luoto (2019) uses just fourth moments).

Definition 1 (Co-skewness and co-kurtosis). Let $\hat{\varepsilon}_t(\tilde{\mathbf{A}}) = \tilde{\mathbf{A}}^{-1}\mathbf{u}_t$. Define *co-skewness* to be:

$$CS_{ijk} \equiv \mathbb{E}\left[\frac{\mathbf{e}_i'\hat{\varepsilon}_t(\tilde{\mathbf{A}})}{\sigma(\mathbf{e}_i'\varepsilon_t(\tilde{\mathbf{A}}))} \frac{\mathbf{e}_j'\hat{\varepsilon}_t(\tilde{\mathbf{A}})}{\sigma(\mathbf{e}_j'\varepsilon_t(\tilde{\mathbf{A}}))} \frac{\mathbf{e}_k'\hat{\varepsilon}_t(\tilde{\mathbf{A}})}{\sigma(\mathbf{e}_k'\varepsilon_t(\tilde{\mathbf{A}}))} \right], \tag{2.6}$$

where $\mathbf{e}_i$ is a selection vector and $\sigma(.)$ denotes the standard deviation. Similarly, define the *co-kurtosis* as:

$$CK_{ijk\ell} \equiv \mathbb{E}\left[\frac{\mathbf{e}_i'\hat{\varepsilon}_t(\tilde{\mathbf{A}})}{\sigma(\mathbf{e}_i'\varepsilon_t(\tilde{\mathbf{A}}))} \frac{\mathbf{e}_j'\hat{\varepsilon}_t(\tilde{\mathbf{A}})}{\sigma(\mathbf{e}_j'\varepsilon_t(\tilde{\mathbf{A}}))} \frac{\mathbf{e}_k'\hat{\varepsilon}_t(\tilde{\mathbf{A}})}{\sigma(\mathbf{e}_k'\varepsilon_t(\tilde{\mathbf{A}}))} \frac{\mathbf{e}_\ell'\hat{\varepsilon}_t(\tilde{\mathbf{A}})}{\sigma(\mathbf{e}_\ell'\varepsilon_t(\tilde{\mathbf{A}}))} \right]. \tag{2.7}$$

While the correlation is also a natural co-moment to consider, it is uninformative about the true rotation matrix. This underlies the lack of identification in models that use only second moments, including Gaussian VARs.

Remark 1 (Zero cross-correlation). It is straightforward to see that mechanically, the cross-correlations are zero for any rotation matrix $\mathbf{Q}$. Cross-correlations are proportional to the off-diagonal elements of the covariance matrix, given by:

$$\text{Var}[\tilde{\mathbf{A}}^{-1}\mathbf{u}_t] = \mathbf{Q}'(\mathbf{\Sigma}^{\text{tr}})^{-1}\text{Var}[\mathbf{u}_t](\mathbf{\Sigma}^{\text{tr}'})^{-1}\mathbf{Q} = \mathbf{Q}'(\mathbf{\Sigma}^{\text{tr}})^{-1}\mathbf{\Sigma}(\mathbf{\Sigma}^{\text{tr}'})^{-1}\mathbf{Q} = \mathbf{Q}'\mathbf{Q} = \mathbf{I}_p. \tag{2.8}$$

Since all off-diagonal elements of the covariance matrix are zero, so are the cross-correlations.

Remark 2 (Constant co-kurtosis under Gaussianity.). For any $\mathbf{Q}$, the co-kurtosis under Gaussianity is equal to 1 if i, j, k, ℓ consist of two distinct pairs, it is equal to 3 if $i = j = k = \ell$, and otherwise it is equal to zero. Isserlis' Theorem (Isserlis, 1918) implies that the co-kurtosis under Gaussianity is given by:

$$CK_{ijk\ell} = CR_{ij}CR_{k\ell} + CR_{jk}CR_{i\ell} + CR_{ki}CR_{j\ell}, \tag{2.9}$$

where CR_{ij} denotes the correlation between forecast errors i and j. For $i = j = k = \ell$, the correlations are all unity and the formula yields the kurtosis of the standard normal variable of three. From Remark 1, we know that all cross-correlations are zero. Thus, if i, j, k, ℓ consist of two pairs, then the right hand side of (2.9) is equal to 1. In all other cases, all terms in (2.9) are zero.

Indeed, one can generalize the last statement for the Gaussian case:

Remark 3 (Identical distribution in Normal Case). Note that the distribution of estimated shocks $\hat{\varepsilon}_t$ is $\mathcal{N}(\mathbf{0}, \mathbf{I}_p)$, independent of $\mathbf{Q}$. This follows directly from the fact that $\mathbf{u}_t$ is mean zero and from equation (2.8). Since $\hat{\varepsilon}_t \overset{d}{=} \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$ffo, it follows that all moments are independent of $\mathbf{Q}$. Thus, they are uninformative.

The main result here is Proposition 1. It shows that under the true $\mathbf{Q}$, the higher co-moments are the same as those of the Gaussian distribution.

Lemma 1 (Independence implies zero co-moments of structural shocks). Let $i(\ell) \in \{1, \ldots, n\}, \ell = 1, \ldots, m$ for $m \leq \bar{m} + 1$. (a) If m is *odd*, then the mth moment $\mathbb{E}[\prod_{\ell=1}^{m} \epsilon_{i(\ell)}]$ is zero unless all $i(\ell)$ are equal. (b) If m is *even*, then the mth moment is zero if there is an ℓ such that $i(\ell) \neq i(\ell') \forall \ell' \neq \ell$.

Proof. Suppose that one of the $i(\ell)$ is different from the others, and without loss of generality let $i(m) \neq i(\ell) \forall \ell < m$. Then

$$\mathbb{E}\left[\prod_{\ell=1}^{m} \epsilon_{i(\ell)}\right] = \mathbb{E}\left[\prod_{\ell=1}^{m-1} \epsilon_{i(\ell)}\right] \mathbb{E}[\epsilon_{i(m)}] = \mathbb{E}\left[\prod_{\ell=1}^{m-1} \epsilon_{i(\ell)}\right] \times 0 = 0,$$

where the first equality follows from independence, the second equality from the distributional assumptions, and the last equality from the fact that $m - 1 \leq \bar{m}$ and the first $\bar{m}$ moments are finite. $\square$

Proposition 1 (The "right" $\mathbf{Q}$ implies the same co-moments as Gaussian distribution). Under Assumption 1, if $\tilde{\mathbf{A}} = \mathbf{A}$, then if one i, j, k (and ℓ, if applicable) is distinct the co-skewness is zero and the co-kurtosis is zero unless the indices occur in pairs, in which case it is unity. These moments are the same as for the Gaussian case.

Proof. Consider the case of the co-kurtosis. The proof for the co-skewness is a simplified version of the same.

For future reference let

$$\hat{\mathbf{A}} \equiv \tilde{\mathbf{A}}^{-1}\mathbf{A} = [\hat{a}_{ij}]_{i,j=1}^p. \tag{2.10}$$

(Note that $\hat{\mathbf{A}}$ is a rotation matrix, i.e., $\hat{\mathbf{A}} \in \mathcal{O}^p$: $\hat{\mathbf{A}}\hat{\mathbf{A}}' = \mathbf{Q}'\mathbf{\Sigma}^{\mathrm{tr}\,-1}\mathbf{A}\mathbf{A}'(\mathbf{\Sigma}^{\mathrm{tr}\,\prime})^{-1}\mathbf{Q} = \mathbf{Q}'\mathbf{\Sigma}^{\mathrm{tr}\,-1}\mathbf{\Sigma}(\mathbf{\Sigma}^{\mathrm{tr}\,\prime})^{-1}\mathbf{Q} = \mathbf{Q}'\mathbf{Q} = \mathbf{I}_p$.)

Note that $\hat{\boldsymbol{\varepsilon}}_t(\tilde{\mathbf{A}}) = \tilde{\mathbf{A}}^{-1}\mathbf{A}\boldsymbol{\varepsilon}_t = \hat{\mathbf{A}}\boldsymbol{\varepsilon}_t$. Using this fact and the fact that all standard deviations of $\hat{\boldsymbol{\varepsilon}}_t$ are unity from equation (2.1) implies from (2.7) that the co-kurtosis is proportional to:

$$\begin{aligned}
CK_{ijk\ell} &= \mathbb{E}\left[\mathbf{e}_i'\hat{\boldsymbol{\varepsilon}}_t(\tilde{\mathbf{A}})\mathbf{e}_j'\hat{\boldsymbol{\varepsilon}}_t(\tilde{\mathbf{A}})\mathbf{e}_k'\hat{\boldsymbol{\varepsilon}}_t(\tilde{\mathbf{A}})\mathbf{e}_\ell'\hat{\boldsymbol{\varepsilon}}_t(\tilde{\mathbf{A}})\right] \\
&= \mathbb{E}[\hat{\mathbf{A}}_{i\circ}\boldsymbol{\varepsilon}_t\hat{\mathbf{A}}_{j\circ}\boldsymbol{\varepsilon}_t\hat{\mathbf{A}}_{k\circ}\boldsymbol{\varepsilon}_t\hat{\mathbf{A}}_{\ell\circ}\boldsymbol{\varepsilon}_t] \\
&= \sum_{\iota_i=1}^p \hat{a}_{i,\iota_i}\mathbb{E}[\epsilon_{\iota_i,t}\hat{\mathbf{A}}_{j\circ}\boldsymbol{\varepsilon}_t\hat{\mathbf{A}}_{k\circ}\boldsymbol{\varepsilon}_t\hat{\mathbf{A}}_{\ell\circ}\boldsymbol{\varepsilon}_t] \\
&= \sum_{\iota_i=1}^p \hat{a}_{i,\iota_i} \sum_{\iota_j=1}^p \hat{a}_{j,\iota_j} \sum_{\iota_k=1}^p \hat{a}_{k,\iota_k} \sum_{\iota_\ell=1}^p \hat{a}_{\ell,\iota_\ell}\mathbb{E}[\epsilon_{\iota_i,t}\epsilon_{\iota_j,t}\epsilon_{\iota_k,t}\epsilon_{\iota_\ell,t}].
\end{aligned} \tag{2.11}$$

Now, consider $\tilde{\mathbf{A}} = \mathbf{A}$. Then $\hat{\mathbf{A}} = \mathbf{I}_p$, so that $\hat{a}_{i,\iota_i} = \mathbf{1}_{i=\iota_i}$ etc. Plugging in:

$$CK_{ijk\ell} \propto \mathbb{E}[\epsilon_{i,t}\epsilon_{j,t}\epsilon_{k,t}\epsilon_{\ell,t}].$$

Since $i < j, k, \ell$, it follows from Lemma 1 that the co-kurtosis is zero.

Consider the four indices forming two pairs, so that $\mathbb{E}[\epsilon_{i,t}\epsilon_{j,t}\epsilon_{k,t}\epsilon_{\ell,t}] = \mathbb{E}[\epsilon_{\iota,t}^2\epsilon_{\kappa,t}^2]$ for some $\iota \neq \kappa$. Because of independence, we then have that:

$$\mathbb{E}[\epsilon_{\iota,t}^2\epsilon_{\kappa,t}^2] = \mathbb{E}[\epsilon_{\iota,t}^2]\mathbb{E}[\epsilon_{\kappa,t}^2] = 1.$$

Comparing the results here with Remark 2 shows that the co-kurtosis is the same as in the Gaussian case.

For the co-skewness, we have by analogy that for $\tilde{\mathbf{A}} = \mathbf{A}$:

$$CS_{ijk} \propto \mathbb{E}[\epsilon_{i,t}\epsilon_{j,t}\epsilon_{k,t}].$$

Now, unless $i = j = k$, and in particular of $i < j, k$, the co-skewness is zero. Since the Gaussian distribution is symmetric with zero skewness, the co-skewness is also the same as in the Gaussian case. $\square$

Remark 4 (Symmetry implies zero co-skewness). If the distributions of all $\epsilon_{i,t}$ are also symmetric around zero, then the co-skewness is always zero.

Imitating the steps in the proof of Proposition 1 yields:

$$CS_{ij} = \sum_{\iota_i=1}^{p} \hat{a}_{i,\iota_i} \sum_{\iota_j=1}^{p} \hat{a}_{j,\iota_j} \sum_{\iota_k=1}^{p} \hat{a}_{k,\iota_k} \mathbb{E}[\epsilon_{\iota_i,t}\epsilon_{\iota_j,t}\epsilon_{\iota_k,t}].$$

If $\iota_i \neq \iota_j$ or $\iota_i \neq \iota_k$, it follows from Lemma 1 that the expectation term is zero. This leaves $\mathbb{E}[\epsilon_{\iota_i,t}^3]$. Under symmetry, $\mathbb{E}[\epsilon_{\iota_i,t}^3] = \mathbb{E}[(-\epsilon_{\iota_i,t})^3] = -\mathbb{E}[\epsilon_{\iota_i,t}^3]$ and, therefore, also equals zero.

Proposition 1 suggests natural tests of whether identification schemes satisfy our independence assumptions.

Define $\boldsymbol{\varepsilon}_t(\mathbf{Q}, \boldsymbol{\phi})$ as $\mathbf{Q}^{-1}(\boldsymbol{\Sigma}^{\mathrm{tr}})^{-1}\mathbf{u}_t$ where $\boldsymbol{\phi}$ is the vector of reduced form parameters, $\mathbf{u}_t$ is the reduced form error corresponding to the parameters in $\boldsymbol{\phi}$, $\boldsymbol{\Sigma}^{\mathrm{tr}}$ is the Cholesky factor of $\boldsymbol{\Sigma}$ corresponding to the parameters in $\boldsymbol{\phi}$ and $\mathbf{Q}$ is a candidate orthonormal matrix.

For the moment-based test, we define a column vector $f(\mathbf{y}_t, \mathbf{Q}, \boldsymbol{\phi})$ consisting of all $\begin{pmatrix} p+2 \\ 3 \end{pmatrix} - p$ elements of the form:

$$\mathbf{e}_i'\boldsymbol{\varepsilon}_t(\mathbf{Q}, \boldsymbol{\phi})\mathbf{e}_j'\boldsymbol{\varepsilon}_t(\mathbf{Q}, \boldsymbol{\phi})\mathbf{e}_k'\boldsymbol{\varepsilon}_t(\mathbf{Q}, \boldsymbol{\phi})$$

such that i, j, $k = 1, ...p$ and $|i, j, k| > 1$ along with all $\begin{pmatrix} p+3 \\ 4 \end{pmatrix} - p$ elements of the form

$$\mathbf{e}_i'\boldsymbol{\varepsilon}_t(\mathbf{Q}, \phi)\mathbf{e}_j'\boldsymbol{\varepsilon}_t(\mathbf{Q}, \phi)\mathbf{e}_k'\boldsymbol{\varepsilon}_t(\mathbf{Q}, \phi)\mathbf{e}_\ell'\boldsymbol{\varepsilon}_t(\mathbf{Q}, \phi) - \mathbf{1}\{i, j, k, \ell \text{ form two distinct pairs})$$

such that i, j, k, $\ell = 1, ...p$ and $|i, j, k, \ell| > 1$. Here $|\{\}|$ denotes the cardinality of the set. From Proposition 1, under Assumption 1, we then have the moment conditions $Ef(\mathbf{y}_t, \mathbf{Q}, \boldsymbol{\phi}) = 0$ for the true matrix $\mathbf{Q}$. Given $\hat{\boldsymbol{\phi}}$, the usual estimate of the reduced form parameters, we can consider the GMM test statistic:

$$S(\mathbf{Q}) = \Sigma_{t=1}^T f(\mathbf{y}_t, \mathbf{Q}, \hat{\boldsymbol{\phi}})[\Sigma_{t=1}^T f(\mathbf{y}_t, \mathbf{Q}, \hat{\boldsymbol{\phi}})f(\mathbf{y}_t, \mathbf{Q}, \hat{\boldsymbol{\phi}})']^{-1}\Sigma_{t=1}^T f(\mathbf{y}_t, \mathbf{Q}, \hat{\boldsymbol{\phi}}).$$

If this were evaluated at the true reduced form parameters, then subject to sufficient moments existing, $TS(Q)$ would have a χ^2 null asymptotic distribution. But evaluating it at the estimated parameters introduces an extra source of randomness (see Appendix A). We thus resort to a bootstrap to simulate its null distribution instead. The bootstrap is implemented as follows. For each hypothesized $\mathbf{Q}$ we take the structural residuals $\boldsymbol{\varepsilon}_t(\mathbf{Q}, \hat{\boldsymbol{\phi}})$. We sample this vector with replacement, and do so for each element separately so that we are imposing independence in the bootstrap sample, as in (2.4). Then, given $\mathbf{Q}$ and $\hat{\boldsymbol{\phi}}$ we construct the bootstrap sample, re-estimate $\boldsymbol{\phi}$ and compute a bootstrap version of $S(\mathbf{Q})$. The exercise is computationally quite intensive because a new bootstrap is required for each candidate $\mathbf{Q}$. Naturally, we reject any $\mathbf{Q}$ where the observed test statistic lies in the upper tail of the bootstrap distribution. The method is the S-set method of Stock and Wright (2000), with ϕ treated as a strongly identified parameter and $\mathbf{Q}$ as a weakly identified parameter for which we wish to conduct inference.

8

The GMM test is based on an assumption of independence (as in Assumption 1), and rules out stochastic volatility. For example, assume that $\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \omega_t^2 \mathbf{I}_p), \mathbb{E}[\omega_t^2] = 1$. In that case, $\mathbb{E}[\epsilon_1^2 \epsilon_2^2] = \mathbb{E}[\mathbb{E}[\epsilon_1^2 \epsilon_2^2 | \omega^2] \omega^4] = \mathbb{E}[\omega^4] \neq 1$, unless ω is degenerate. However, the GMM test can still work with conditional independence as in Assumption 2, as long as we drop the $p(p-1)/2$ symmetric kurtosis conditions:

$$\mathbb{E}[\epsilon_{\iota,t}^2 \epsilon_{\kappa,t}^2] = 1, \iota < \kappa \leq m.$$

We henceforth refer to the GMM tests for independence with and without the symmetric kurtosis conditions as the GMM1 and GMM2 tests, respectively. While the GMM2 test is robust to stochastic volatility, we are not exploiting any time-varying volatility for identification. That would involve considering other moments, like cross-products of squared shocks and lagged squared shocks, but we leave that for future work.

2.2.2 Non-parametric test of independence

While the moment-based criteria have intuitive appeal, co-skewness or co-kurtosis are just two particular ways of looking at the data and may fail to detect dependence. This is most apparent in the case of symmetric distributions and co-skewness, but may also be the case with co-kurtosis. We therefore turn now to a direct test of the independence assumption (2.4a), following Hoeffding (1948) and Blum et al. (1961). We henceforth refer to this henceforth as the BKR test. This approach has the advantage that, in population, the test statistic is zero iff the joint distribution is the product of independent marginals.

Starting from (2.4a), let $F(\boldsymbol{\varepsilon})$ denote the joint distribution of the shocks $\boldsymbol{\varepsilon}$ and compare it to the product measure of its marginals. Call the resulting difference D:

$$D(\boldsymbol{\varepsilon}) \equiv F(\boldsymbol{\varepsilon}) - \prod_{j=1}^p F_j(\epsilon_j), \tag{2.12}$$

where $F_j(\epsilon_j) = F([\mathbf{1}_{(j-1)'} \times \infty, \epsilon_j, \mathbf{1}(p-j)' \times \infty])$ is the jth marginal distribution of F.

Define the population expectation of the Hoeffding (1948)-Blum et al. (1961) statistic as the integral over D with respect to the joint distribution[2]

$$\Delta(F) \equiv \int D(\boldsymbol{\varepsilon})^2 dF(\boldsymbol{\varepsilon}) = \mathbb{E}[D(\boldsymbol{\varepsilon})]. \tag{2.13}$$

Lemma 2 (Hoeffding (1948) Theorem 3.1 extended to $p \geq 2$)**.** When $F(x_i, \ldots, x_p)$ is a continuous distribution function with continuous joint and marginal probability densities, then $\Delta(F) = 0$ if and only if the components of $\mathbf{x}$ are mutually independent, i.e., $F(\mathbf{x}) = \prod_{j=1}^p F_j(x_j)$.

Proof. (This proof closely follows Hoeffding (1948), who considered $p = 2$.)
$\Rightarrow$: $F(\mathbf{x}) = \prod_{j=1}^p F_j(x_j)$ implies that $D(\mathbf{x}) = 0$ and thus that $\Delta(F) = 0$.

---

[2]Note that Bergsma and Dassios (2014) define the Blum et al. (1961) statistic with respect to the product measure of the marginals.

$\Leftarrow$: Suppose that $F(\mathbf{x}) \neq \prod_{j=1}^{p} F_j(x_j)$ so that $D(\mathbf{x}) \neq 0$. Given our assumption on F, the function $d(\mathbf{x}) = f(\mathbf{x}) - \prod_j f_j(x_j)$ is continuous. We have that:

$$D(\mathbf{x}) = \int_{-\infty}^{x_1} \ldots \int_{-\infty}^{x_p} d(u_1, \ldots, u_p) du_1 \ldots du_p.$$

$D(\mathbf{x}) \neq 0$ thus implies that $d(\mathbf{x}) \neq 0$ and because

$$\int_{-\infty}^{\infty} \ldots \int_{-\infty}^{\infty} d(u_1, \ldots, u_p) du_1 \ldots du_p = \int_{-\infty}^{\infty} \ldots \int_{-\infty}^{\infty} dF(\mathbf{u}) - \prod_{j=1}^{p} \int_{-\infty}^{\infty} dF_j(u_j) = 1 - 1 = 0,$$

there exists a rectangle A in $\mathbb{R}^p$ such that $d(\mathbf{x}) > 0$ for $\mathbf{x} \in A$. Thus $D(\mathbf{x}) \neq 0$ almost everywhere in A and, because $d(\mathbf{x}) = f(\mathbf{x}) - \prod_j f_j(x_j) > 0$ and $f_j(x_j) \geq 0$ it follows that $f(\mathbf{x}) > 0$ in A. Thus:

$$\Delta(F) \geq \int_A D(\mathbf{x})^2 f(\mathbf{x}) d\mathbf{x} > 0.$$

Thus $D(\mathbf{x}) \neq 0$ implies that $\Delta(F) > 0$. $\qquad\square$

Proposition 2 (The "right" $\mathbf{Q}$ implies a zero population test statistic)**.** Let $F(\boldsymbol{\varepsilon} | \tilde{\mathbf{A}})$ be the distribution of the VAR shocks $\boldsymbol{\varepsilon}$ given the impact matrix $\tilde{\mathbf{A}}$. $\Delta(F(\circ | \tilde{\mathbf{A}}) = 0$ if and only if $\tilde{\mathbf{A}}$ renders the distribution of $\boldsymbol{\varepsilon}$ independent. In particular, if $\tilde{\mathbf{A}} = \mathbf{A}$ then $\Delta(F(\circ | \mathbf{A})) = 0$.

Under Gaussianity, we have $\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$ and thus $\Delta(F(\circ | \tilde{\mathbf{A}})) = 0$ for all $\tilde{\mathbf{A}}$.

Remark 5. Continuity matters. If we allow discrete distributions, then Hoeffding (1948) already contains a counterexample where the test statistic is zero despite dependence. This can be circumvented by integrating not with respect to dF_{12} but with respect to the product measure of the marginals. Sometimes (e.g., Bergsma and Dassios (2014, p. 1008)), this is how BKR is written.

Turning to the sample version of the independence test, let $\hat{F}_{\mathbf{Q},\phi}(x_1, \ldots x_p)$ be the empirical joint cdf of all p elements of $\boldsymbol{\varepsilon}_t(\mathbf{Q}, \phi)$ and let $\hat{F}_{\mathbf{Q},\phi}(x_i)$ be the empirical marginal cdf of $e_i' \boldsymbol{\varepsilon}_t(\mathbf{Q}, \phi)$. The BKR test statistic is:

$$BKR(\mathbf{Q}, \phi) = \int (\hat{F}_{\mathbf{Q},\phi}(x_1, \ldots x_p) - \Pi_{i=1}^{p} \hat{F}_{\mathbf{Q},\phi}(x_i))^2 d\hat{F}_{\mathbf{Q},\phi}(x_1, \ldots x_p).$$

The null hypothesis is of independence, as in Assumption 1. The test statistic evaluated at the true $\mathbf{Q}$ and ϕ has an asymptotic null distribution that is derived by BKR, but we are going to instead test a hypothesized value of $\mathbf{Q}$ replacing ϕ with its estimate $\hat{\phi}$. We use the same bootstrap algorithm used for evaluating the GMM critical values to simulate the null distribution of $BKR(\mathbf{Q}, \hat{\phi})$.

Given $\hat{\phi}$, we thus have a set of permissible rotations of the Cholesky factor consistent with the

data. Using the GMM and BKR tests, these sets are given by:

$$\mathbb{Q}_{\alpha,S} = \{\mathbf{Q} \in \mathcal{O}^p | S(Q) \leq \widehat{CV}_{\alpha,S}\}, \tag{2.14a}$$

$$\mathbb{Q}_{\alpha,BKR} = \{\mathbf{Q} \in \mathcal{O}^p | BKR(Q) \leq \widehat{CV}_{\alpha,BKR}\}, \tag{2.14b}$$

where $\widehat{CV}_{\alpha,\circ}$ denote the bootstrapped critical values.

2.2.3 Testing point-identification

Point-identified VARs – or other time series models, such as the state space representation of linear DSGE models or factor models – imply a certain rotation $\hat{\mathbf{Q}}$. We can test the assumptions underlying these models by checking whether $\hat{\mathbf{Q}} \in \mathbb{Q}_{\alpha,S}$ or $\hat{\mathbf{Q}} \in \mathbb{Q}_{\alpha,BKR}$.

Note that in many applications, not all p shocks in VARs are identified. In that case, we only test for independence of the $\tilde{p}$ shocks, where $1 < \tilde{p} \leq p$. In other cases, researchers may also want to focus on a subset of shocks that are of particular interest, in order to test them with higher power.

2.2.4 Refining set identification

Our overall approach to inference is frequentist and incorporates sign restrictions. It makes sense to view our test as a refinement of inference via sign restrictions because even in the case where independent components analysis achieves point identification, it does so only up to a reordering of the columns of $\mathbf{Q}$ and flipping the signs of each column. Thus the researcher must have some beliefs about the sign of impulse responses to be able to exploit the implication of independence.

Granziera et al. (2018) use a grid to form a $100(1 - \alpha_1)\%$ confidence set for $\mathbf{q}_1$, the first column of $\mathbf{Q}$ and then for each selected $\mathbf{q}_1$ they form a $100(1 - \alpha_2)\%$ for each impulse response to the first structural shock, taking into account the uncertainty about ϕ. Since ϕ is root-T consistently estimable, the latter is a standard inference problem. They then take the union of all of these confidence sets, and by appeal to the Bonferroni inequality, this has a coverage of $100(1 - \alpha_1 - \alpha_2)\%$. Although Granziera et al. (2018) test a hypothesis for a single column of $\mathbf{Q}$ rather than for the whole matrix, and correspondingly are concerned with identification of a single shock, their approach can also be used to test hypotheses on the entire matrix $\mathbf{Q}$.

Our approach also takes a grid of values of $\mathbf{Q}$. For each of these, we test if the first column, $\mathbf{q}_1$, satisfies the sign restrictions using a test of size α_3 and the test of Granziera et al. (2018).[3] For those $\mathbf{Q}$s that are accepted, we then proceed to test if $\mathbf{Q}$ satisfies independence of the structural errors, using either the GMM or BKR test statistic and we use a test of size α_4. Finally, conditional on any $\mathbf{Q}$ that satisfies the first two conditions, we form a $100(1- \alpha_5)\%$ confidence interval for each impulse response to the first structural shock. We take the convex hull of all of these confidence intervals, and by appeal to the Bonferroni inequality, this has coverage of $100(1 - \alpha_3 - \alpha_4 - \alpha_5)\%$. To preserve

---

[3]Alternatively, we could test if the entire matrix $\mathbf{Q}$ satisfies the sign restrictions, but in this paper we are focusing on identifying a single shock.

the possible non-convexity of the confidence intervals, we take the union of the confidence intervals conditional on $\mathbf{Q}$, rather than the minimum and maximum of the bounds of those confidence interval as in Granziera et al. (2018).

We also consider a method that treats the structural VAR as point identified, assuming non-normality and in the spirit of the existing literature using higher order moments for identification. The idea is to take the moment conditions $E(\boldsymbol{\varepsilon}_t \otimes \mathbf{y}_{t-l}) = 0$ for $l = 1, ...k$, $E(vech(\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t' - I)) = 0$ and $Ef(\mathbf{y}_t, \mathbf{Q}, \boldsymbol{\phi}) = 0$ and use these as an overidentified GMM problem to get a point estimate of $\boldsymbol{\phi}$ and $\mathbf{Q}$ yielding point estimates of the impulse responses along with conventional standard errors. This is essentially the GMM approach of Lanne and Luoto (2019). But it is subject to a lack of identification in the case of normal errors, and possibly weak identification even with non-normal errors. Because the $\mathbf{Q}$ matrix is only identified in this way up to a permutation of the columns and flipping the signs within each column, it seems reasonable to impose that the chosen column of the estimated $\mathbf{Q}$, in conjunction with the estimated $\boldsymbol{\phi}$, delivers instantaneous impulse response functions that meet the sign restriction.

The presence of higher moments or a non-parametric dependence test makes the Bonferroni approach more attractive here than the Delta method (Gafarov et al., 2018) or a projection approach (Gafarov et al., 2016). The extra constraints coming from the requirement that the higher order dependence be small prevent the quasi-analytic characterization of the bounds on IRFs needed to apply the Delta method. For the projection method, there are two challenges: First, when combining the moment conditions that test for higher order dependence into a joint Wald ellipsoid with all moments identifying the reduced form parameters, the presence of the potentially imprecisely estimated higher moments can make the Wald ellipsoid very conservative when projected onto the IRF space. While we could limit this problem by using a Bonferroni approach for combining first and second moments with higher order moments, inverting the Wald ellipsoid presents a second, numerical challenge: The presence of the higher order moments makes it hard to invert the Wald ellipsoid. It is also unclear how to invert the Wald ellipsoid to yield a possibly non-convex confidence interval.

2.3 Statistical identification and non-convex confidence intervals

Lemma 3. Identification through independence is only local: If $\mathbf{Q}$ satisfies the independence criterion in sample, then there exists at least one $\tilde{\mathbf{Q}} \neq \mathbf{Q}$ that is also an orthonormal matrix that also satisfies the independence criterion in sample. This $\tilde{\mathbf{Q}}$ is outside a neighborhood of $\mathbf{Q}$.

Proof. Suppose that $\mathbf{Q} = [\mathbf{q}_1, \ldots, \mathbf{q}_p]$ satisfies the independence test of choice. Then any genuine reordering $\tilde{\mathbf{Q}} = [\mathbf{q}_{i1}, \ldots, \mathbf{q}_{in}] \neq \mathbf{Q}$, where $\{i_1, \ldots, i_p\} = \{1, \ldots, p\}$, also satisfies the independence criterion. Since the independence tests are invariant to the ordering of shocks, the test statistics are unchanged. Note that this $\tilde{\mathbf{Q}}$ is not local, i.e., it lies outside an open set around $\mathbf{Q}$ in $\mathcal{O}_p$. $\square$

In population, any $\mathbf{Q}$ that just transforms the marginals, i.e., any combination of reordering or sign flips, is a non-local transform. It thus continues to satisfy the independence criteria. For the

GMM statistic, neither reordering nor sign flips has any effect on the sample statistic. For the BKR statistic, reordering has no effect on the sample statistic, but sign flips change the empirical value. In any event identification is only local, and this is why we impose the independence requirement as a refinement of sign restrictions.

Bacchiocchi and Kitagawa (2020) analyze a similar failure of global identification despite local identification that can arise from zero restrictions.

Because of the local statistical identification, the refined confidence interval resultant from our proposed procedure may be non-convex. This contrasts with sign restrictions, which, by themselves, lead to convex confidence intervals under mild conditions (see Granziera et al., 2018). Consider, for example, a bivariate VAR with positively correlated forecast errors. If the sign restrictions require comovement or the reduced form covariance is low, the sign-identified set for $\mathbf{q}$, relative to the lower Cholesky factor, admits both a $\mathbf{q}$ with very large or very small responses, as well as the intermediate responses. The statistically identified set consists of four regions, each convex, whose centers are 90 degrees apart. When the sign-identified set is wide enough and the true $\mathbf{q}$ close enough to one of its boundaries, then not just one, but two of the statistically identified sets intersect with the sign-identified set and the resulting confidence interval is non-convex.[4]

3   Monte Carlo Simulations

Applying the proposed estimator to a simulated data-generating process (DGP) with Gaussian and non-Gaussian shocks serves as a proof of concept: It verifies that with Gaussian shocks, there is no material set refinement. In contrast, with heavy tailed and especially with skewed shocks, the estimated identified set shrinks considerably.

We consider a bivariate VAR:

$$
\mathbf{y}_t = \left(\begin{array}{cc} 0.97 & 0 \\ 0.5 & 0.9 \end{array} \right) \mathbf{y}_{t-1} + A\boldsymbol{\varepsilon}_t,
$$

where $A = \boldsymbol{\Sigma}^{tr}\mathbf{Q}$, $\boldsymbol{\Sigma} = \left(\begin{array}{cc} 1 & 0.3 \\ 0.3 & 2 \end{array} \right)$ and $\mathbf{Q}$ is the orthonormal matrix $\left(\begin{array}{cc} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{array} \right)$ and we specify that $\theta = \frac{1}{3}\pi$. This means that a positive value of the first structural shock increases the first variable and reduces the second. It also implies that the true parameter is in the interior of the identified set; below we also consider θ near the boundary of the identified set. We consider five different distributional specifications for $\boldsymbol{\varepsilon}_t$: (i) normally distributed, (ii) t distributed on 4 degrees of freedom, (iii) one shock is normal and the other is t distributed on 4 degrees of freedom, (iv) skewed t and (v) Laplace. The first three are all symmetric; the last two add in skewness. In each case we consider four confidence intervals: (i) impulse responses by the method of Granziera et al. (2018) with $\alpha_1 = \alpha_2 = 0.05$, so a nominal 90% confidence interval, (ii) impulse responses by

---

[4]Below, we provide an empirical example of this case. In principle, the reduced form parameter uncertainty is too large; however, the confidence interval for the IRFs can be convex even if the confidence interval for $\mathbf{q}$ is not.

our proposed method using a GMM test with $\alpha_3 = \alpha_4 = .025$ and $\alpha_5 = 0.05$, so again a nominal 90% confidence interval, (iii) the same thing but with the BKR test for independence instead, and (iv) a GMM confidence interval assuming identification, with nominal 90% coverage. The GMM confidence intervals use the GMM1 test, including symmetric kurtosis terms. Results for the GMM2 test are similar, but are not shown to keep the results easier to read. The coverage and mean width of these different intervals are reported in Figures 1-5 in a sample size of $T = 500$.

Figure 1: Simulation Results with Normal Errors

This figure reports the coverage and average length of impulse responses in the simulated VAR for each horizon in the Monte Carlo simulation described in the text. Blue Line: Sign Restrictions, Black Line: GMM, Red Line: BKR, Green Line: GMM based on assumption of identification. The sample size is $T = 500$.

First, Figure 1 deals with the normal case. The sign restrictions method and the proposed refinements have coverage that is a bit above the nominal level, and all have essentially identical width, because the independence restriction has no bite in this case. The GMM estimate assuming point identification gives much shorter confidence intervals. But, as is to be expected given the lack of identification, their coverage is way too low. In this figure, the sign restrictions only,

14

Figure 2: Simulation Results with t(4) errors
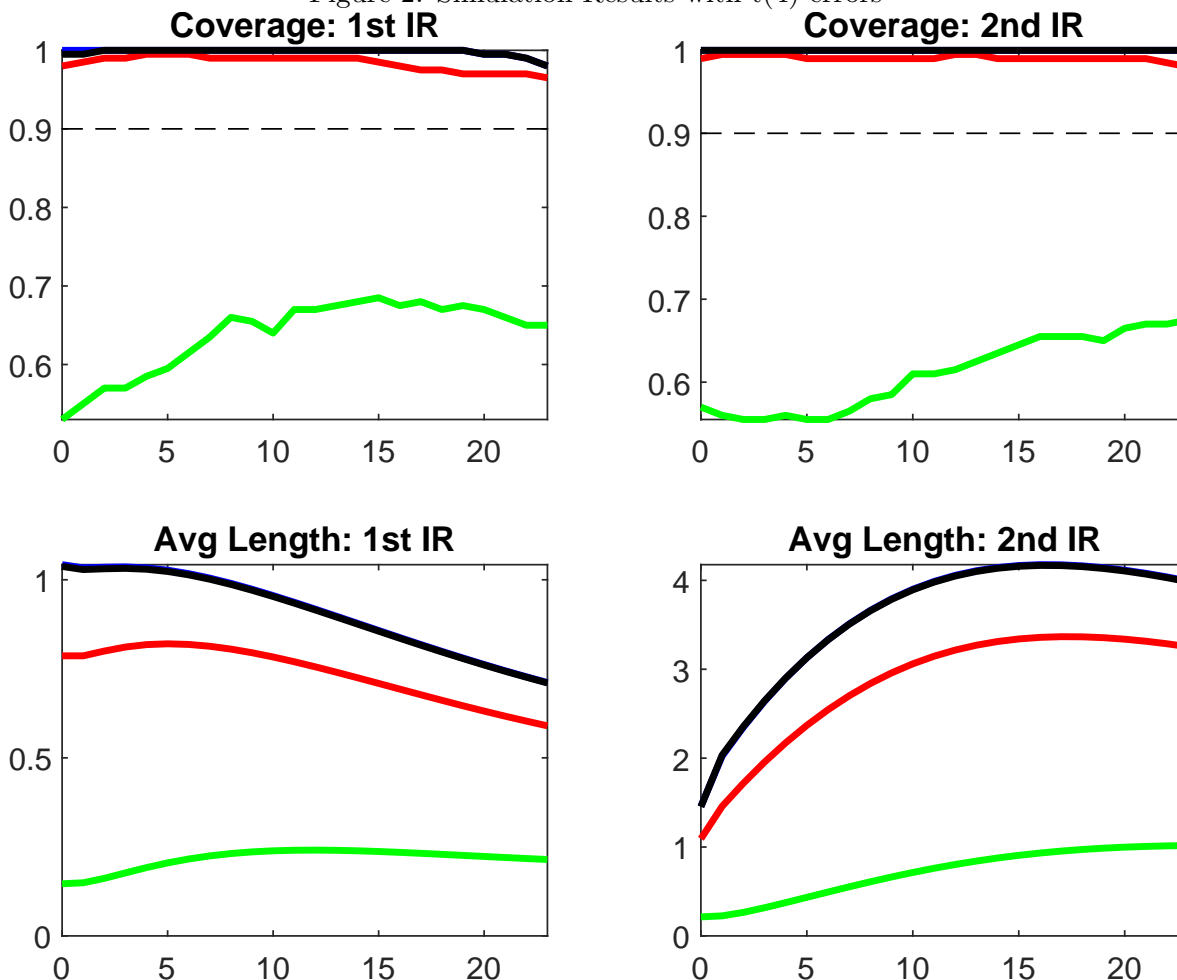
This figure reports the coverage and average length of impulse responses in the simulated VAR for each horizon in the Monte Carlo simulation with described in the text. Blue Line: Sign Restrictions, Black Line: GMM, Red Line: BKR, Green Line: GMM based on assumption of identification. The sample size is $T = 500$.

GMM refinement and BKR refinement yield results that are so close that the three lines are indistinguishable.

Next, Figure 2 deals with the t(4) case. The GMM estimate assuming point identification has coverage that is better but still too low at around 60 percent. The BKR refinement reduces width relative to the sign identified intervals. The improvement is small, but noticeable. The GMM refinement reduces width too, but the reduction is so small that the sign restrictions and GMM refinement results are indistinguishable in these figures.

Data generating processes with different shocks or skewness yield larger effects of the refinement. Figure 3 shows this by combining one normal shock with one t(4) shock. GMM assuming identification does no better than in the pure t(4) case. However, the BKR refinement reduces the average width of the confidence interval by half now. The GMM refinement still has no visible effect.

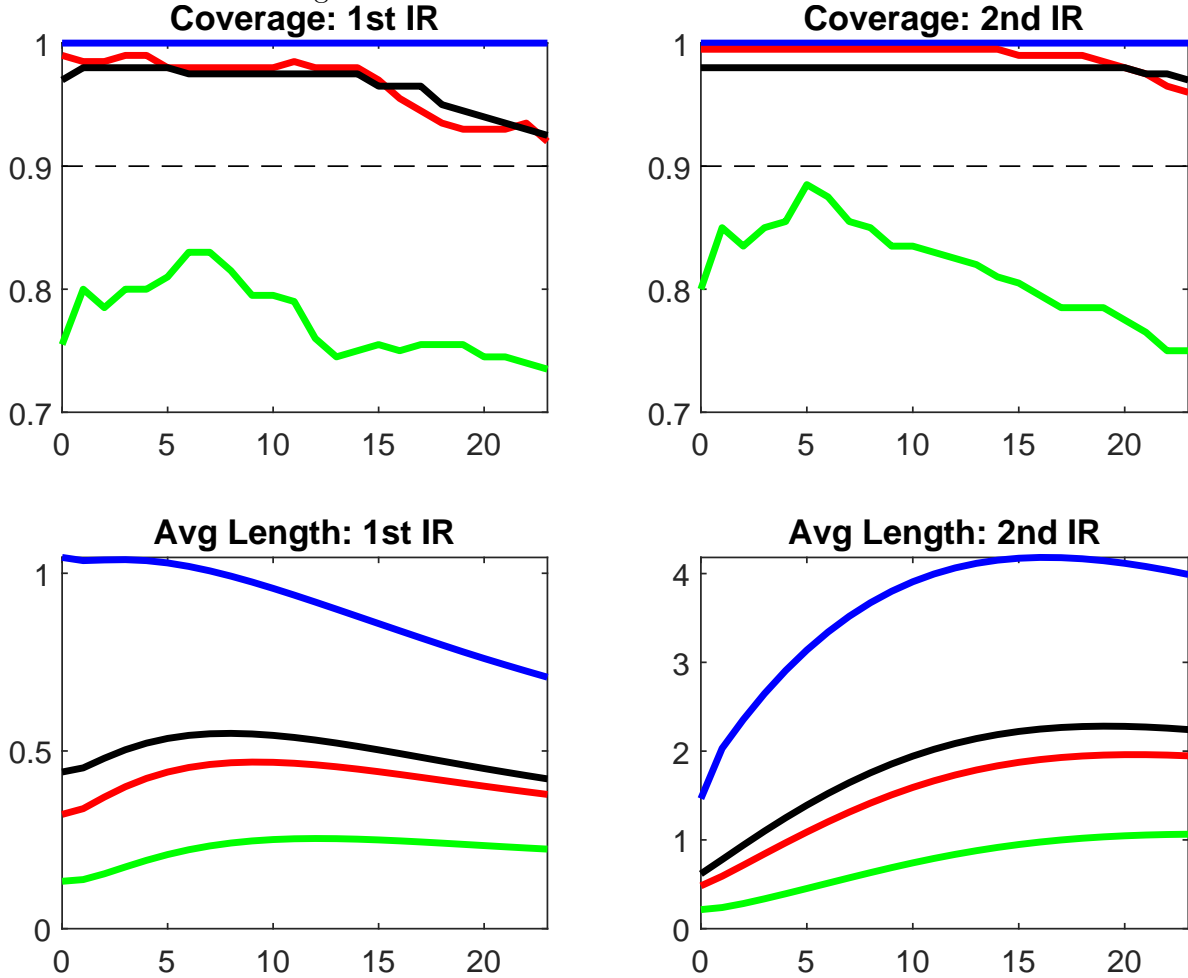Figure 3: Simulation Results with one Normal Error and one t(4) Error
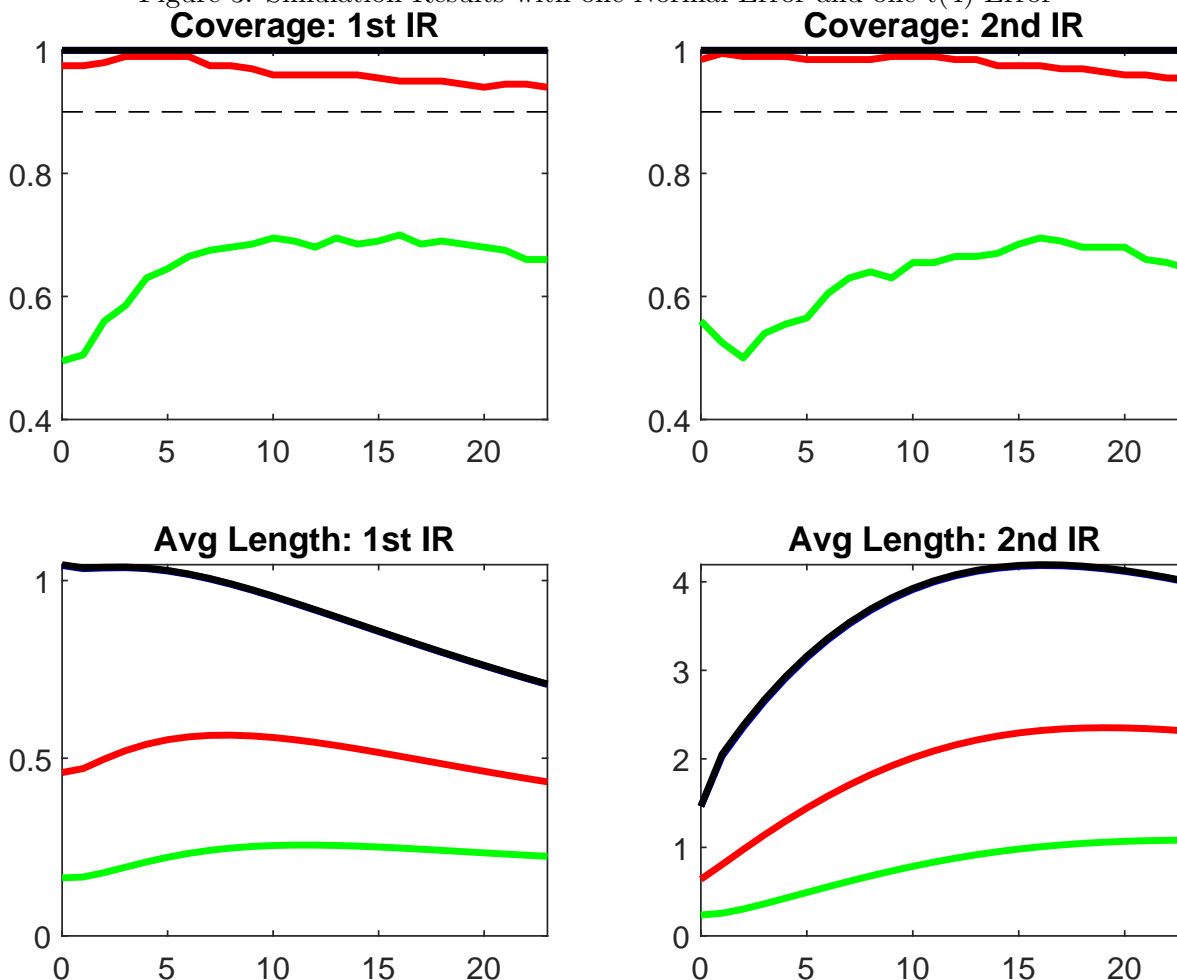
This figure reports the coverage and average length of impulse responses in the simulated VAR for each horizon in the Monte Carlo simulation with described in the text. Blue Line: Sign Restrictions, Black Line: GMM, Red Line: BKR, Green Line: GMM based on assumption of identification. The sample size is $T = 500$.

Figure 4 concerns the asymmetric Laplace distribution with skewness of 2 and excess kurtosis of six. The sign restrictions method and the proposed refinements once again control coverage. The refinements very substantially reduce the width of the confidence intervals, especially in the case of the BKR test. The GMM estimate assuming point identification again gives shorter confidence intervals, but doesn't really control coverage.

Figure 5 deals with the skew t case (Hansen, 1994). The distribution has eight degrees of freedom and a skewness parameter of 0.5. In this case, the sign restrictions method and the proposed refinements continue to control coverage. The refinements noticeably reduce the width of the confidence intervals, especially in the case of the BKR test. In contrast, with a plain t(8) distribution without skewness, the refinement has barely any effect (not shown). The GMM estimate assuming point identification gives shorter confidence intervals, but the effective coverage is still a bit below the nominal level.

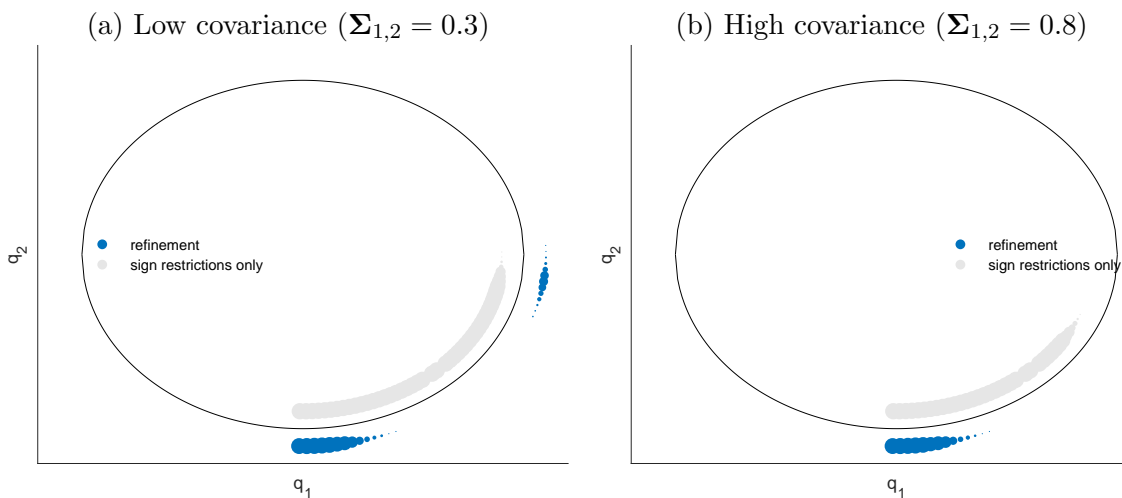Figure 4: Simulation Results with Laplace Errors

This figure reports the coverage and average length of impulse responses in the simulated VAR for each horizon in the Monte Carlo simulation with described in the text. Blue Line: Sign Restrictions, Black Line: GMM, Red Line: BKR, Green Line: GMM based on assumption of identification. The sample size is $T = 500$.

Next, we consider two variations: A higher covariance (increasing $\mathbf{\Sigma}_{1,2}$ from 0.3 to 0.8), and setting the true $\mathbf{q}$ near the boundary—being $[\cos(\theta), -\sin(\theta)]'$ with $\theta = 0.49\pi$. Increasing the covariance between the reduced form forecast errors shrinks the identified set when identifying a shock with negative comovement between the two IRFs.[5] Figure 6 plots the identified set for $\mathbf{q}$ based on simulations for the skew t(8) distribution with the true $\mathbf{q}$ near the boundary. With a low covariance between forecast errors, the sign restrictions imply an identified set for $\mathbf{q}$ of almost 90 degrees (gray shaded area). Adding in the independence refinement based on the BKR test yields both a neighborhood around the true $\mathbf{q}$ at the southern pole and some elements of the neighborhood around its 90-degree rotation (the blue area). The most extreme members of the sign-identified set for $\mathbf{q}$ survive in many simulations and there is consequently little gain in inference about the

---

[5]See, e.g., Amir-Ahmadi and Drautzburg (2021) for a discussion of how covariance and restrictions determine the identified set.

Figure 5: Simulation Results with Skew t Errors

This figure reports the coverage and average length of impulse responses in the simulated VAR for each horizon in the Monte Carlo simulation with described in the text. Blue Line: Sign Restrictions, Black Line: GMM, Red Line: BKR, Green Line: GMM based on assumption of identification. The sample size is $T = 500$.

convex hull of the confidence interval. This illustrates that the independence refinement does not always give a convex confidence interval, and may be most helpful in conjunction with both sign and magnitude restrictions—a point to which we shall return in the empirical work below. In contrast, when the covariance between the forecast errors is higher, the sign restrictions imply sharper identification and the independence refinement gives a further reduction in the size of the confidence set to a small convex set.

Figure 7 compares the average length of the confidence intervals with sign restrictions and with the BKR refinement for four cases: The true **q** in the interior with high or low covariance, and the true **q** on the boundary with high or low covariance. The location of the truth makes no difference for the width of the confidence intervals with sign restrictions, but a higher covariance narrows the average width slightly. Relative to using sign restrictions alone, the statistical refinement helps in all cases, but by different amounts in different cases. With the interior truth, the width does not

Figure 6: Simulated Identified Set for **q** with Skew t Errors. Low vs High covariance. Truth near boundary.

(a) Low covariance ($\mathbf{\Sigma}_{1,2} = 0.3$) (b) High covariance ($\mathbf{\Sigma}_{1,2} = 0.8$)

This figure reports the support on the identified set across simulations. The wider the shaded areas, the larger the fraction of simulations that include this region in the identified set. True $\mathbf{q}_0 = [0.0314, -0.9995] = [\cos(\theta), -\sin(\theta)]'$ with $\theta = \frac{1}{3}\pi$. The sample size is $T = 500$.

Figure 7: Simulation Results with Skew t Errors: Comparison of high and low covariance and truth in interior vs boundary.
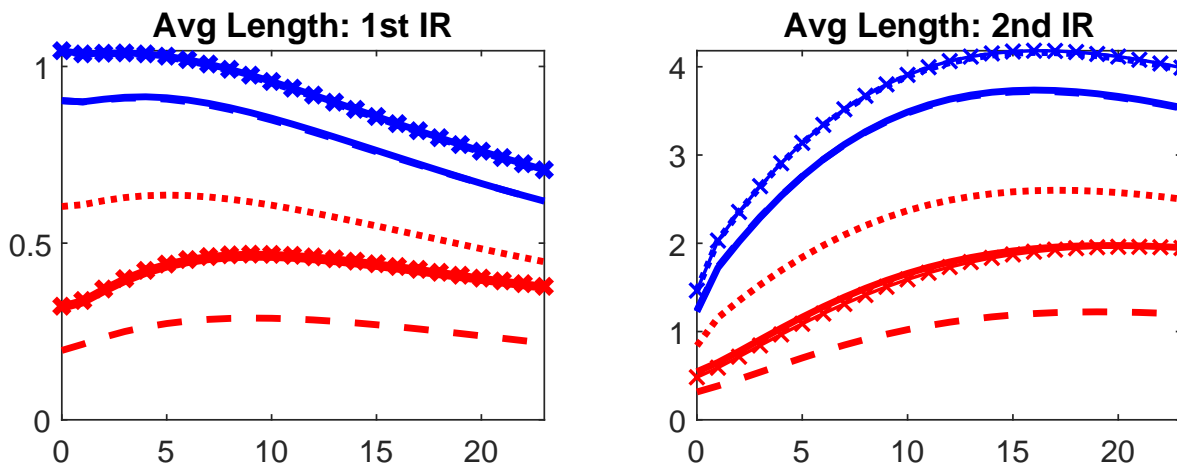
This figure reports the average length of impulse responses in the simulated VAR for each horizon in the Monte Carlo simulation with described in the text. Blue Line: Sign Restrictions, Red Line: BKR. Solid line: interior, high covariance; Dashed line: boundary, high covariance; Dotted line: boundary, low covariance; Solid line with markers: interior, low covariance. Interior: $\mathbf{q}_0 = [0.5, -0.8660] = [\cos(\theta), -\sin(\theta)]'$ with $\theta = \frac{1}{3}\pi$. Boundary: $\mathbf{q}_0 = [0.0314, -0.9995] = [\cos(\theta), -\sin(\theta)]'$ with $\theta = 0.49\pi$. The sample size is $T = 500$.

depend on the covariance, as the more informative sign restrictions only confirm the information provided by the statistical identification. In contrast, with the true **q** near the boundary, the low covariance leads to less precise inference relative to the interior truth. The opposite is true with

19

high covariance and the average width with the refinement is about one third of the average width with sign restrictions only.[6]

Additional simulations suggest that many non-normal distributions require large samples to yield statistical identification, and even with a very long time series, the statistical identification yields sets of non-negligible measure. In the Appendix (Figure A2), we consider sample sizes of $T \in \{500, 1000, 2500, 5000\}$ and assess the power to reject independence of *iid* shocks, which are drawn either from the t_4 distribution or the skewed t_8 distribution. The true **q** is the south pole of the unit circle. For the t_4 distribution, the BKR test needs sample sizes of 2500 for the statistical identification to reduce the identified set to segments near the poles of the unit circle. The GMM test requires a sample size of 5000, and even then the statistical identification has no effect in about 10% of the simulations. For the skewed t_8 distribution, in contrast, both tests and all sample sizes yield reliable identification across all simulations. But even with a sample size of 5000, the statistically identified set does not resemble a collection of singletons.

In summary, in these simulations, the sign restrictions method and proposed refinements are the only methods that control coverage. The refinements are of no use in the normal case, but do little harm in this case. The BKR test generally has more power than the GMM test. Meanwhile, the refinements can help, but the improvement is only material in the case of errors with different DGPs, skewed errors or for unusually long samples. The improvement is more pronounced when the sign restrictions yield an identified set for **q** that spans less than 90 degrees and when the true **q** is near the boundary of this identified set.

4 Empirics

4.1 Testing point identification

One use of our proposed method is to test a specific claimed identification scheme. We demonstrate this with three empirical examples.

Bloom (2009): Uncertainty shocks. We take the "kitchen sink" (full specification) VAR model of stock market valuations, uncertainty shocks, and six macroeconomic control variables from Bloom (2009). We follow Bloom in using 12 lags. Since Bloom (2009) only takes a stand on the Cholesky ordering of the stock index and the uncertainty measure, we only test the independence of these two shocks.

Both identified shocks exhibit heavy and asymmetric tails, as we show in Appendix C.[7] The departure from the normal distribution is pronounced, with visible left skewness of shocks to stock returns and right skewness to uncertainty shocks.

---

[6]With the truth on the boundary, our proposed procedure has actual coverage equal to the nominal coverage of 0.90 for the parameter that is on the boundary. Otherwise, it continues to be conservative. See Figure A1.

[7]Figure A3 plots the residuals and the BKR statistic over a rolling 10-year window, and Figure A4 shows QQ plots relative to the standard normal distribution. The spiked nature of the uncertainty residual is because the underlying variable has a large point mass at zero – and is otherwise continuously distributed.

Table 1 shows the test statistic for the BKR and GMM test. We can reject the independence of these shocks according to the BKR test and the GMM2 test at the 95% level, and according to the GMM1 test at the 90% level. Thus the identification scheme is clearly rejected.

Test	Test statistic	95% CV	90% CV
BKR	0.83	0.05	0.04
GMM1	31.25	36.7	22.78
GMM2	21.23	9.28	7.74

Critical values obtained via bootstrap of slope coefficients, intercept, and Cholesky factor.

Table 1: Test of independence of uncertainty shock and stock market valuation shock in Bloom (2009) VAR, 1963m1 to 2008m12.

Blanchard and Quah (1989): Demand and supply shocks. Blanchard and Quah (1989) have pioneered the identification of TFP shocks using the long-run restriction that only TFP shocks affect deviations of GDP from its trend growth in the long run. In their bivariate system, that restriction is sufficient to also identify the other shocks, dubbed a demand shock.

Here, we use the replication data for Blanchard and Quah (1989) provided in Herwartz (2019). The data cover 1948q2 through 1987q4.[8] The data have been demeaned and detrended prior to estimation. Following Blanchard and Quah (1989) and Herwartz (2019), we estimate the VAR with eight lags.

The identified shocks have moderately heavier tails than the normal distribution according to our estimates, as we show in Appendix D.[9] Table 2 implies that the identifying assumption clearly fails the independence test according to the BKR test, while the GMM tests reject independence more narrowly.

Test	Test statistic	95% CV	90% CV
BKR	0.0480	0.0340	0.0295
GMM1	14.79	15.01	11.31
GMM2	7.29	6.22	5.21

Critical values obtained via bootstrap of slope coefficients, intercept, and impact matrix.

Table 2: Test of independence of demand and supply shock in Blanchard and Quah (1989)-style VAR, 1948q2 to 1987q4.

Herwartz (2019) considers a statistical point identification in the original Blanchard and Quah (1989) sample. He finds that the original identification scheme is only weakly rejected by the data, with a *p*-value of 0.15. This indicates that our dependence criteria are more sensitive toward

---

[8]Similar results hold for various VAR specifications with updated data vintages and a sample period from 1948q1 through 2019q4.

[9]Figure A5 plots the residuals and the BKR statistic over a rolling 10-year window, and Figure A6 shows QQ plots relative to the standard normal distribution.

detecting dependence than the criterion used in Herwartz (2019). Rejection of the identified shocks as dependent, however, need not imply substantively different conclusions. The statistical point identification in Herwartz (2019), which maximizes the p-value of an independence test, yields IRFs that are very close to those in Blanchard and Quah (1989).

Smets and Wouters (2007): DSGE model. Smets and Wouters (2007) present a DSGE model for the U.S. economy, whose forecast performance is comparable to that of reduced form Bayesian VARs. Unlike those reduced form models, it allows for structural analysis.

Here, we take this DSGE model as an example of a state space model. It models seven time series from 1966q1 to 2004q4 in the main Smets and Wouters (2007) sample, and yields seven identified shocks. We focus on the posterior mode, which, as T goes to infinity, converges to the maximum likelihood estimator according to the Bernstein-von-Mises Theorem. We use the estimated DGP to generate 500 bootstrap samples. We test the joint independence of all shocks, but also test the independence of government spending shocks and monetary policy shocks. This is motivated by the finding of monetary accommodation of fiscal shocks in Drautzburg (2020).

Most of the seven identified shocks exhibit heavier tails than the normal distribution; see Appendix E.[10] The wage markup shock and the investment shock are the exception. In the Appendix figures and in the analysis below, we standardize the shocks to have unit standard deviation. However, unlike in identified VAR models, the shocks are generally correlated in sample. To isolate higher moment dependence, we thus consider "pre-whitened" structural shocks by multiplying the shocks by the Cholesky factor of the in-sample covariance matrix.

Table 3 shows the results of the independence test. The non-parametric BKR test fails to reject the joint independence of all shocks. However, when focusing on the monetary and fiscal policy shocks, the BKR test rejects independence. For the GMM statistics, since $p = 7$, the number of higher order moments is high. The critical values are two orders of magnitude above the test statistic when testing all shocks – a clear failure to reject independence. This may reflect the difficulty in estimating so many higher moments.

Test	All shocks			Fiscal and Monetary policy shocks		
	Test statistic	95% CV	90% CV	Test statistic	95% CV	90% CV
BKR	0.0070	0.0206	0.0186	0.0599	0.0352	0.0311
GMM1	2.7×10^5	3.1×10^7	1.4×10^7	4.1	22.1	16.4
GMM2	1.0×10^4	1.3×10^4	1.1×10^4	4.0	7.1	6.0

Critical values obtained via bootstrap with 500 replications.

Table 3: Test of independence of identified shocks in Smets and Wouters (2007).

---

[10]Figure A7 plots the residuals and the BKR statistic over a rolling 10-year window, and Figure A8 shows QQ plots relative to the standard normal distribution.

4.2 Refining set identification: Labor market application

In their influential work on sign-restrictions, Baumeister and Hamilton (2015) consider a bivariate VAR that characterizes the U.S. labor market. The VAR has overall employment growth and the growth of compensation per hour, and is estimated on quarterly data with eight lags.

Here, we revisit their application. To obtain a longer, monthly sample, we analyze the manufacturing labor market, rather than the overall labor market.[11] We include manufacturing employment and the real average hourly earnings in the manufacturing sector. While the data are available starting in 1939m1, we use the sample from 1945m1 to 2019m12, to exclude the effects of WWII and the COVID-19 pandemic.[12] We use 24 lags and an intercept in the estimation.

Figure 8 shows the OLS residuals in the top two panels. The top panel shows the log real earnings residual from 1947m1 to 2019m12, normalized by its standard deviation as the thick, blue line. The thick, blue line in the middle panel shows the log employment residual over the same time period. In both panels, the thin, orange line shows the corresponding uncorrelated residuals ("shocks"), computed using a Cholesky factor of the covariance matrix that orders employment first. For earnings, the Cholesky-identified shock by assumption equals the normalized residual. For employment, the two are allowed to differ, but happen to largely coincide.

Visual inspection of the residuals suggests departures from normality. The employment residual (and Cholesky-identified shock) exceeds five standard deviations in absolute value eight times, all between 1947 and 1971; under the standard normal distribution, the probability of having a five standard deviation shock is less than one in a million. The earnings residual has relatively thinner tails, but still has heavier tails than the standard normal distribution. The Q-Q plots in Figure A10 in the Appendix, which plot quantiles of residuals against quantiles from the standard normal distribution, make the same point more explicitly. Since departure from Normality is a necessary condition to get a bite from the refinement we propose, this is reassuring.

However, in our framework it is dependence of otherwise uncorrelated shocks that refines the identification. Heavy tails of the marginal distributions need not translate to such dependence. Indeed, when we turn to the data, we observe that dependence is driven by the part of the sample. The bottom panel in Figure 8 computes the BKR statistic over rolling ten-year windows. Over the first forty years of the sample, the BKR test statistic is roughly centered around its 5% critical value (in the absence of estimation uncertainty), which is shown as the dashed line. The rolling windows that end in the late 1990s detect much higher dependence of the two Cholesky-identified shocks. The refinement of identification comes from ruling out **Q**-rotations of the Cholesky-identified shocks that also exhibit excessive dependence, using an appropriate critical value.

We next turn to constructing confidence sets for IRFs by the Bonferroni method discussed in section 2.2.4. Following Baumeister and Hamilton (2015), we identify a negative labor demand shock as a shock that simultaneously lowers employment and earnings. A labor supply shock is

---

[11]Qualitatively, we obtain similar results for the overall labor market, but the results are less pronounced. For the overall labor market, the sample period runs from 1960m1 to 2019m12. See Appendix G for the analogous results.

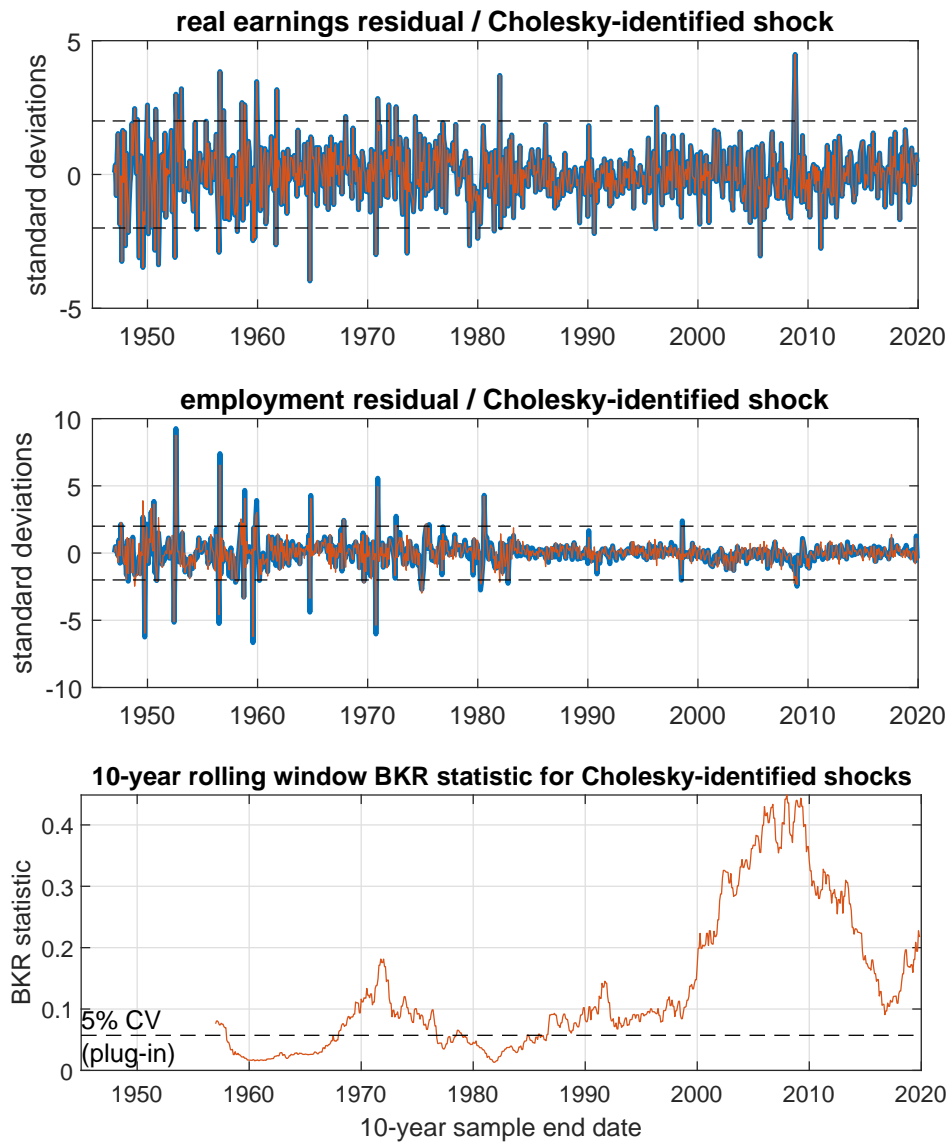[12]See Figure A9 for a time series plot of the raw data and data description.

Figure 8: Real manufacturing earnings residuals (top panel), employment residuals (middle panel), and 10-year rolling window of BKR test of Cholesky-identified shocks (bottom panel): 1947m1–2019m12

identified from wages falling and employment rising. The analysis below imposes these restrictions on impact and two months out, although the results change little with restriction horizons from one to six months.

We first analyze the effects of the pure sign restrictions and the proposed refinement at the MLE, following Moon and Schorfheide (2012). Figure 9 shows the identified set for the first column of $\mathbf{Q}$ on the unit circle as well as the implied impact IRFs for wages and employment on the horizontal and vertical axes, respectively. The results with sign restrictions alone are shown in gray on the interior of the unit circle for the coordinates of the rotation vector. For the labor demand shock, shown on the left, the identified set is a continuous arc from $(0, -1)$ clockwise to about $(-0.95, 0.3)$. This arc therefore includes the negative of either column of the Cholesky-identified shocks. The implied identified set for the impact IRF is shown as the continuous lines parallel to the axes. For the wage response (IRF_1), for example, the identified set is $[0, -0.42]$. For the employment response (IRF_2), the identified set is $[0, -0.46]$.[13]

Figure 9 shows the effects of our proposed refinements based on the inverted Blum et al. (1961) test statistic.[14] The refined identified set for the rotations is shown on the exterior of the unit circle in both red and blue. It is a disjoint set of two arcs that exclude the poles included with sign restrictions alone. The identified set does, however, leave in many rotations on either extreme of the arc identified with sign restrictions alone. Mapped to impact IRFs, the convex hull of the identified set shrinks only marginally compared to sign restrictions alone. The only gain is that now we can bound the effect of a negative demand shock on employment away from zero. The refinement has bite, but this comes mainly from ruling out large parts of the interior of the identified set under sign restrictions alone. These imply that the negative labor demand shock either has a large wage effect and a modest employment effect (blue area), or a large employment effect and a small wage effect (red area).

This reflects the local nature of the statistical identification: If $\mathbf{Q} = [\mathbf{q}_1, \mathbf{q}_2]$ satisfies the independence requirement, then so does $[\mathbf{q}_2, \mathbf{q}_1]$. And, if $\mathbf{q}_2$ in turn satisfies the sign restriction, then it will also form part of the refined identified set. Note that $\pm\mathbf{q}_2$ is orthogonal to $\mathbf{q}_1$, i.e., a ± 90 degree rotation. This gives the potential for the refined identified set to consist of a disjoint set of two arcs, as we see here.

For the labor supply shock, shown to the right, the independence refinement provides a more straightforward tightening of inference. The sign restrictions identify an arc spanning less than 90 degrees and only one statistically identified region survives.

---

[13]Mechanically, the sign restriction on the identified shock requires the two IRFs to move together. Since also the reduced form errors are positively correlated, this restriction has no bite – it is consistent with the identified shock explaining all or nothing of the variation in both variables (see Granziera et al. (2018) and Amir-Ahmadi and Drautzburg (2021)).

[14]Inverting the GMM test statistics has little or no bite in this application: The GMM test statistic that includes symmetric co-kurtosis terms does not reject any $\mathbf{Q}$ on the grid. The GMM test statistic that is consistent with stochastic volatility and omits symmetric co-kurtosis terms, in contrast, rejects about 6% of the grid points. This points to difficulty in estimating these higher moments accurately.
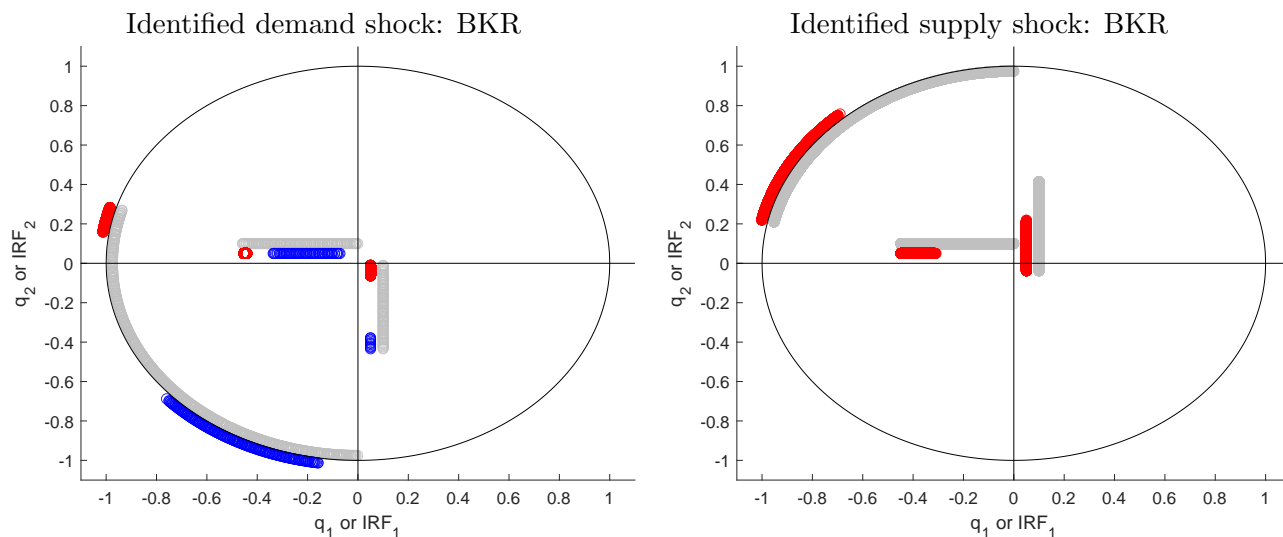
Mapping from $\mathbf{Q}$ to impact IRFs based on Cholesky factor $\widehat{\mathbf{\Sigma}^{\mathrm{tr}}} = \begin{bmatrix} 0.46 & 0 \\ 0.13 & 0.42 \end{bmatrix}$.

Figure 9: Refinement of set identification in Baumeister and Hamilton (2015)-style analysis of manufacturing labor market: Identified set of Cholesky factor rotations and of impact IRFs at the MLE
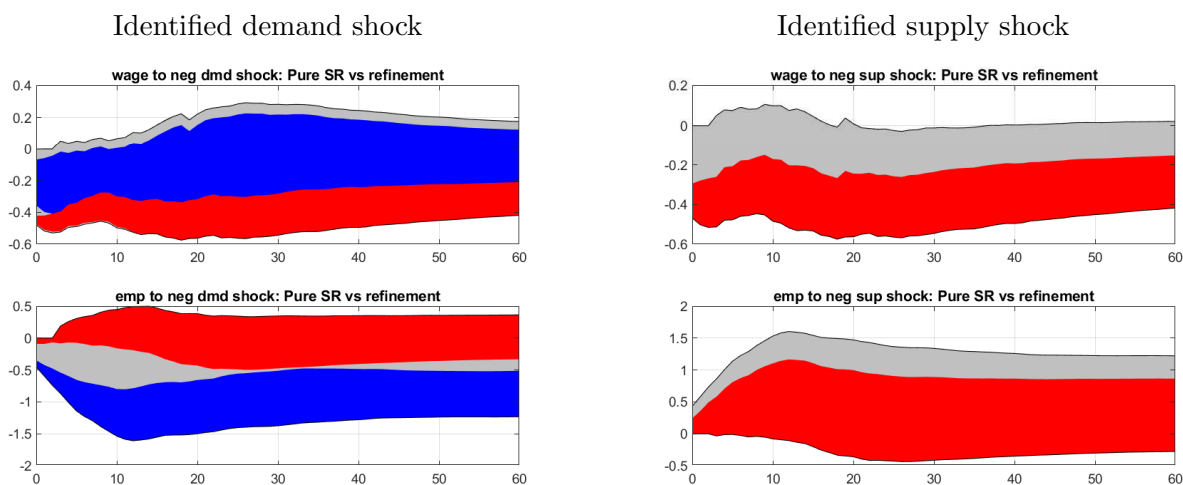

Figure 10: Refinement of set identification in Baumeister and Hamilton (2015)-style analysis of manufacturing labor market: Full confidence interval of IRFs for first 60 months. The areas in grey use sign restrictions alone. The areas shaded in red and blue (for demand shocks) and in red (for supply shocks) incorporate the refinement using the BKR independence test.

The results at the MLE carry over to the full estimation of parameters and rotations. Figure 10 shows the identified sets for IRFs up to five years out and takes into account parameter uncertainty using our full estimation procedure, building on Granziera et al. (2018). The results are similar to

those at the MLE. For the labor demand shock with the sign restrictions only, the only inference one can draw is to bound the magnitude of the employment and wage declines on impact: The 90% confidence intervals are, roughly, $[-0.5, 0]$ – shown as the gray area in Figure 10. Although the convex hull of the refined identified set is not very different from what we get by sign restrictions alone, the confidence set with our refinement is not convex, and instead turns out to consist of two distinct areas: draws with small effects on wages are plotted in red and those with big effects on wages are plotted in blue, as discussed above. Meanwhile, for the labor supply shock, the independence refinement substantially reduces the width of the confidence intervals.

In cases where the independence refinement leaves two distinct arcs in the identified set for a column of **Q**, such as with our identified labor demand shock, it is natural to think of magnitude restrictions, which may give the refinement more bite. Micro evidence suggests that the aggregate labor supply curve is fairly steep (Robert and Mok, 2012) and so it might be reasonable to posit that the labor supply elasticity is no greater than 1. This means that the ratio of the impulse response of a labor demand shock on wages to that on employment must be greater than or equal to 1. Given that the labor demand shock lowers employment, that can in turn be rewritten as a restriction that the effect of the labor demand shock on employment less than on wages must be nonnegative, and that is a sign restriction that can be incorporated in the procedure of Granziera et al. (2018).

Figure 11 shows the identified sets for IRFs to labor demand shocks up to five years out using (i) the sign restrictions and elasticity restrictions imposed at horizons 0, 1 and 2 months alone (in grey) and (ii) also incorporating the independence refinement (in red). Here the independence refinement has a large bite, as it effectively picks one of the two arcs that we found in Figure 9. The labor demand shock has a very small immediate effect on employment and a large effect on wages.
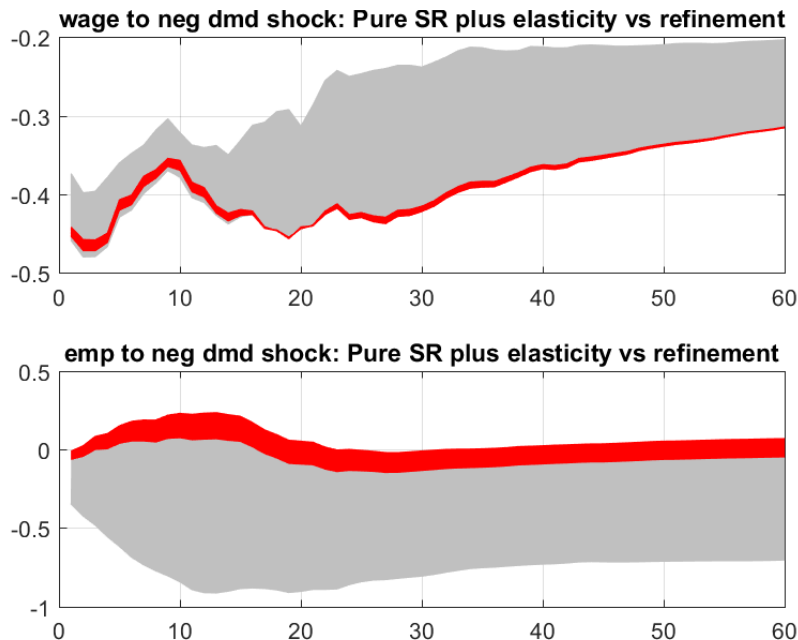
Figure 11: The figure shows the confidence intervals of IRFs to a labor demand shock for the first 60 months. Intervals are shown using sign restrictions and the elasticity restriction that the impact on wages is greater than that on employment (in grey). The red areas add to this the independence refinement based on the BKR test

Similarly, the refined identified set has stark implications for the historical shocks. With the refinement, there are several historical episodes that imply either very large or small to modest labor supply shocks to explain the observed data. Figure 12 lists the ten months in our sample with the largest difference in the absolute labor supply shock between the rotation maximizing and minimizing the BKR statistic. For these ten episodes, Figure 12 shows the identified set of the absolute shock size for the accepted draws in red and blue. For example, in August of 1980 (80m08), the rotations consistent with the inverted BKR test statistic imply a shock size of 2.5 standard deviations or less (blue area), or of around 4.5 standard deviations (red area). The former shock is associated with large employment IRFs to the negative demand shock; the latter with small responses. If narrative shock restrictions (Antolin-Diaz and Rubio-Ramírez (2018), Giacomini et al. (2021)) were available for these episodes, they could be combined to also reduce the convex hull of the confidence set for IRFs, similar to the elasticity bounds.

While the analysis here applies only to employment and wages in the U.S. manufacturing sector, qualitatively similar results hold in the quarterly Baumeister and Hamilton (2015) data for the U.S labor market as a whole, as reported in Appendix G.
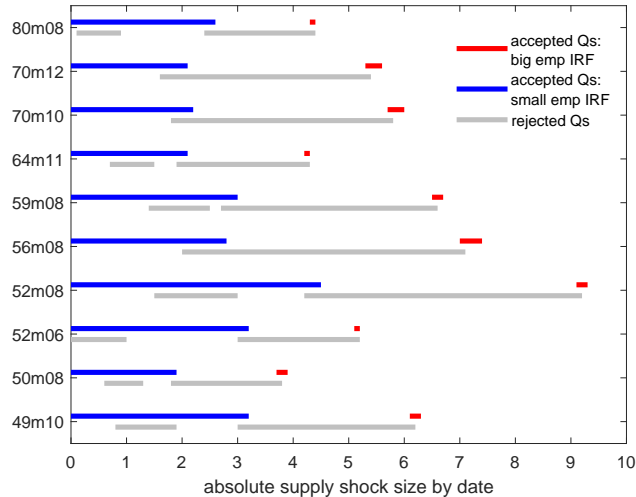
The figure shows the support of the labor supply shock, the shock orthogonal to the identified shock, and how it differs across the accepted and the rejected **Q** draws. For comparison, we plot the absolute shock size, which is measured in standard deviations. The 10 episodes with the largest difference between the **Q**s with the largest and smallest BKR statistics are shown.

Figure 12: Implications of rotations on magnitude of historical supply shocks: Identified set of shock sizes at the MLE

5 Conclusions

We have proposed a method for refining sign identification in vector autoregressions based on ruling out regions in the identified set with significant higher order dependence, such as co-skewness or co-kurtosis. There is a substantial literature on using non-normality for identification of structural VARs, but all of the existing papers assume non-Gaussianity, and thus assume identification. Our focus has been on developing an approach that is robust to weak identification or a complete lack of identification. Besides regularity conditions, the underlying assumption allows for either Gaussian or non-Gaussian shocks. The identification that we obtain is inherently local—for this reason, we combine it with sign restrictions or with both sign and magnitude restrictions. We show that the proposed method controls coverage well in simulations regardless of the distribution of shocks, and it tightens inference when the shocks are heavily skewed. In empirical applications, we have shown that it rules out many rotations and can sharpen inference. Using sign restrictions plus the independence refinement, we find that the confidence intervals for impulse responses for a labor demand shock are non-convex. But adding the independence refinement on top of both sign and magnitude restrictions gives short, and convex, confidence intervals.

References

AMIR-AHMADI, P. AND T. DRAUTZBURG (2021): "Identification and inference with ranking restrictions," Quantitative Economics, 12, 1–39.

ANTOLIN-DIAZ, J. AND J. F. RUBIO-RAMÍREZ (2018): "Narrative Sign Restrictions for SVARs," American Economic Review, 108, 2802–29.

BACCHIOCCHI, E. AND T. KITAGAWA (2020): "Locally- but not globally-identified SVARs," Cemmap Working Paper CWP40/20, Institute for Fiscal Studies.

BAUMEISTER, C. AND J. D. HAMILTON (2015): "Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information," Econometrica, 83, 1963–1999.

BERGSMA, W. AND A. DASSIOS (2014): "A consistent test of independence based on a sign covariance related to Kendall's tau," Bernoulli, 20, 1006–1028.

BERNANKE, B. S. (1986): "Alternative Explanations of the Money Income Correlation," Carnegie Rochester Conference in Public Policy, 25, 49–99.

BLANCHARD, O. J. AND D. QUAH (1989): "The Dynamic Effects of Aggregate Demand and Supply Disturbances," American Economic Review, 79, 655–673.

BLOOM, N. (2009): "The impact of uncertainty shocks," Econometrica, 77, 623–685.

BLUM, J. R., J. KIEFER, AND M. ROSENBLATT (1961): "Distribution Free Tests of Independence Based on the Sample Distribution Function," The Annals of Mathematical Statistics, 32, 485 – 498.

CANOVA, F. AND G. DE NICOLO (2002): "Monetary disturbances matter for business fluctuations in the G-7," Journal of Monetary Economics, 49, 1131–1159.

DRAUTZBURG, T. (2020): "A narrative approach to a fiscal DSGE model," Quantitative Economics, 11, 801–837.

FAUST, J. (1998): "The robustness of identified VAR conclusions about money," Carnegie-Rochester Conference Series on Public Policy, 49, 207–244.

GAFAROV, B., M. MEIER, AND J. MONTIEL OLEA (2016): "Projection Inference For Set-Identified SVARs," Tech. rep., New York University.

——— (2018): "Delta-Method Inference For A Class of Set-Identified SVARs," Journal of Econometrics, 203, 316–327.

GAIGALL, D. (forthcoming): "Hoeffding-Blum-Kiefer-Rosenblatt independence test statistic on partly not identically distributed data," Communications in Statistics - Theory and Methods.

GIACOMINI, R., T. KITAGAWA, AND M. READ (2021): "Identification and Inference Under Narrative Restrictions," Working paper.

GOURIÉROUX, C., A. MONFORT, AND J.-P. RENNE (2017): "Statistical inference for independent component analysis: Application to structural VAR models," Journal of Econometrics, 196, 111 – 126.

GRANZIERA, E., H. MOON, AND F. SCHORFHEIDE (2018): "Inference for VARs Identified with Sign Restrictions," Quantitative Economics.

GUAY, A. (2021): "Identification of structural vector autoregressions through higher unconditional

moments," Journal of Econometrics.

HAMILTON, J. D. (1994): Time series analysis, Princeton university press.

HANSEN, B. E. (1994): "Autoregressive Conditional Density Estimation," International Economic Review, 35, 705–730.

HERWARTZ, H. (2019): "Long-run neutrality of demand shocks: Revisiting Blanchard and Quah (1989) with independent structural shocks," Journal of Applied Econometrics, 34, 811–819.

HOEFFDING, W. (1948): "A Class of Statistics with Asymptotically Normal Distribution," The Annals of Mathematical Statistics, 19, 293 – 325.

HYVÄRINEN, A., J. KARHUNEN, AND E. OJA (2004): Independent Component Analysis, Adaptive and Cognitive Dynamic Systems: Signal Processing, Learning, Communications and Control, Wiley.

ISSERLIS, L. (1918): "On a Formula for the Product-Moment Coefficient of any Order of a Normal Frequency Distribution in any Number of Variables," Biometrika, 12, 134–139.

KILIAN, L. AND H. LÜTKEPOHL (2018): Structural Vector Autoregressive Analysis, no. 9781107196575 in Cambridge Books, Cambridge University Press.

KILIAN, L. AND D. P. MURPHY (2012): "Why Agnostic Sign Restrictions Are Not Enough: Understanding The Dynamics of Oil Market VAR Models," Journal of the European Economic Association, 10, 1166–1188.

LANNE, M. AND J. LUOTO (2019): "GMM Estimation of Non-Gaussian Structural Vector Autoregression," Journal of Business & Economic Statistics, 0, 1–13.

——— (2020): "Identification of Economic Shocks by Inequality Constraints in Bayesian Structural Vector Autoregression," Oxford Bulletin of Economics and Statistics, 82, 425–452.

LANNE, M., M. MEITZ, AND P. SAIKKONEN (2017): "Identification and estimation of non-Gaussian structural vector autoregressions," Journal of Econometrics, 196, 288–304.

LEWIS, D. J. (2021): "Identifying shocks via time-varying volatility," Review of Economic Studies.

MONTIEL OLEA, J. L., M. PLAGBORG-MØLLER, AND E. QIAN (2021): "SVAR Identification From Higher Moments: Has the Simultaneous Causality Problem Been Solved?" Working Paper.

MOON, H. R. AND F. SCHORFHEIDE (2012): "Bayesian and Frequentist Inference in Partially Identified Models," Econometrica, 80, 755–782.

RIGOBON, R. (2003): "Identification through heteroskedasticity," Review of Economics and Statistics, 85, 777–792.

ROBERT, R. AND S. MOK (2012): "A Review of Recent Research on Labor Supply Elasticities," Congressional Budget Office Working Paper 2012-12.

SIMS, C. A. (1980): "Comparison of Interwar and Postwar Business Cycles: Monetarism Reconsidered," American Economic Review, 70, 250–57.

SMETS, F. AND R. WOUTERS (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," American Economic Review, 97, 586–606.

STOCK, J. H. AND J. H. WRIGHT (2000): "GMM With Weak Identification," Econometrica, 68, 1055–1096.

UHLIG, H. (1998): "The robustness of identified VAR conclusions about money : A comment," Carnegie-Rochester Conference Series on Public Policy, 49, 245–263.

——— (2003): "What moves GNP?" draft, Humboldt Universität zu Berlin.

——— (2005): "What are the effects of monetary policy on output? Results from an agnostic identification procedure," Journal of Monetary Economics, 52, 381–419.

Appendices

A Parameter dependence of moment-based test

Let $\mathbf{Y} = [\mathbf{y}_1', \ldots, \mathbf{y}_T']'$ be the $T \times p$ vector of observables. Let $\mathbf{U}$ be the corresponding vector of residuals, with $\hat{\mathbf{U}}$ its sample analogue. $\mathbf{E} = [\boldsymbol{\varepsilon}_1', \ldots, \boldsymbol{\varepsilon}_T']'$ are the structural shocks: $\mathbf{E} = \mathbf{BU}$. Let $\mathbf{Y} = \mathbf{XA} + \mathbf{B}^{-1}\mathbf{E}$. Then:

$$
\begin{aligned}
[\hat{\boldsymbol{\varepsilon}}_{\circ 1}, \ldots, \hat{\boldsymbol{\varepsilon}}_{\circ p}] = \hat{\mathbf{E}} &= (\mathbf{Y} - \mathbf{X}\hat{\mathbf{A}})\hat{\mathbf{B}} \\
&= (\mathbf{Y} - \mathbf{X}(\mathbf{A}_0 + \hat{\mathbf{A}} - \mathbf{A}_0))\hat{\mathbf{B}} \\
&= (\mathbf{Y} - \mathbf{X}\mathbf{A}_0)(\mathbf{B}_0 + \hat{\mathbf{B}} - \mathbf{B}_0) - \mathbf{X}(\hat{\mathbf{A}} - \mathbf{A}_0))\hat{\mathbf{B}} \equiv \mathbf{E} + \mathbf{E}\mathbf{B}^{-1}(\hat{\mathbf{B}} - \mathbf{B}_0) -
\end{aligned}
$$

Here, $\mathbf{B}^{-1}(\mathbf{B}^{-1})' = \boldsymbol{\Sigma}$. Note that when using the MLE/OLS estimator, we have that (see Hamilton (1994), p. 301):

$$
\sqrt{T}\begin{bmatrix} \text{vec}(\hat{\mathbf{A}} - \mathbf{A}_0) \\ \text{vech}(\hat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}_0) \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \boldsymbol{\Sigma} \otimes \mathbb{E}[\mathbf{X}'\mathbf{X}]^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{V}_{\boldsymbol{\Sigma}} \end{bmatrix}\right),
$$

so that the two parameter vectors are asymptotically independent.

Note that $\hat{\epsilon}_{it} - \epsilon_{it} \xrightarrow{p} 0$ because $\mathbf{u}_t(\hat{\mathbf{B}} - \mathbf{B}_0) \xrightarrow{p} 0$ and $\mathbf{x}_t(\hat{\mathbf{A}} - \mathbf{A}_0)\hat{\mathbf{B}} \xrightarrow{p} \mathbf{x}_t(\hat{\mathbf{A}} - \mathbf{A}_0)\mathbf{B} \xrightarrow{p} 0$.

Now consider moments of the form $\mathbb{E}[\epsilon_{it}^m \epsilon_{jt}^n]$ and their sample analogue:

$$
T^{-1}\sum_{t=1}^{T} \hat{\epsilon}_{it}^m \hat{\epsilon}_{jt}^n = T^{-1}\sum_{t=1}^{T}(\epsilon_{it}^m \epsilon_{jt}^n + m\epsilon_{it}^{m-1}\epsilon_{jt}^n(\hat{\epsilon}_{it} - \epsilon_{it}) + n\epsilon_{it}^m \epsilon_{jt}^{n-1}(\hat{\epsilon}_{jt} - \epsilon_{jt}) + O_p(1/T)).
$$

Consequently:

$$
\begin{aligned}
\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\left(\hat{\epsilon}_{it}^m \hat{\epsilon}_{jt}^n - \epsilon_{it}^m \epsilon_{jt}^n\right) &= \frac{1}{T}\sum_{t=1}^{T}\left(m\epsilon_{it}^{m-1}\epsilon_{jt}^n \sqrt{T}\left(\boldsymbol{\varepsilon}_t'\mathbf{B}^{-1}(\hat{\mathbf{B}} - \mathbf{B}_0) - \mathbf{x}_t'(\hat{\mathbf{A}} - \mathbf{A}_0)\hat{\mathbf{B}}\right)\mathbf{e}_i \right. \\
&\qquad\qquad \left. + n\epsilon_{it}^m \epsilon_{jt}^{n-1}\sqrt{T}\left(\boldsymbol{\varepsilon}_t'\mathbf{B}^{-1}(\hat{\mathbf{B}} - \mathbf{B}_0) - \mathbf{x}_t'(\hat{\mathbf{A}} - \mathbf{A}_0)\hat{\mathbf{B}}\right)\mathbf{e}_j\right) \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad (\text{A.1}) \\
&= \left(\frac{1}{T}\sum_{t=1}^{T} m\epsilon_{it}^{m-1}\epsilon_{jt}^n \boldsymbol{\varepsilon}_t'\right)\mathbf{B}^{-1}\underbrace{\sqrt{T}(\hat{\mathbf{B}} - \mathbf{B}_0)}_{\xrightarrow{d}\mathcal{N}}\mathbf{e}_i
\end{aligned}
$$

$$-\left(\frac{1}{T}\sum_{t=1}^{T}m\epsilon_{it}^{m-1}\epsilon_{jt}^{n}\mathbf{x}_t'\right)\underbrace{\sqrt{T}(\hat{\mathbf{A}}-\mathbf{A}_0)}_{\xrightarrow{d}\mathcal{N}}\underbrace{\hat{\mathbf{B}}}_{\xrightarrow{p}\mathbf{B}_0}\mathbf{e}_i$$

$$+\left(\frac{1}{T}\sum_{t=1}^{T}n\epsilon_{it}^{m}\epsilon_{jt}^{n-1}\boldsymbol{\varepsilon}_t'\right)\mathbf{B}^{-1}\underbrace{\sqrt{T}(\hat{\mathbf{B}}-\mathbf{B}_0)}_{\xrightarrow{d}\mathcal{N}}\mathbf{e}_j$$

$$-\left(\frac{1}{T}\sum_{t=1}^{T}n\epsilon_{it}^{m}\epsilon_{jt}^{n-1}\mathbf{x}_t'\right)\underbrace{\sqrt{T}(\hat{\mathbf{A}}-\mathbf{A}_0)}_{\xrightarrow{d}\mathcal{N}}\underbrace{\hat{\mathbf{B}}}_{\xrightarrow{p}\mathbf{B}_0}\mathbf{e}_j+O_p(1/\sqrt{T})$$

$$(\text{A.2})$$

What does $\frac{1}{T}\sum_{t=1}^{T}n\epsilon_{it}^{m}\epsilon_{jt}^{n-1}\mathbf{x}_t'$ converge to? Note that $x_t=[\mathbf{y}_{t-1},\ldots,\mathbf{y}_{t-\ell}]$. Given the DGP, we thus have that $\boldsymbol{\varepsilon}_t\perp\!\!\!\perp\mathbf{x}_t$. Also we use that for $i\neq j$.

$$\frac{1}{T}\sum_{t=1}^{T}n\epsilon_{it}^{m}\epsilon_{jt}^{n-1}\mathbf{x}_t'\xrightarrow{p}\mathbb{E}[n\epsilon_{it}^{m}\epsilon_{jt}^{n-1}\mathbf{x}_t']=\mathbb{E}[n\mathbb{E}[\epsilon_{jt}^{n-1}|\mathbf{x}_t,\epsilon_{it}^{m}]\epsilon_{it}^{m}\mathbf{x}_t']=\mathbb{E}[n\mathbb{E}[\epsilon_{jt}^{n-1}]\epsilon_{it}^{m}\mathbf{x}_t']$$

Now consider the following cases:

$n=1$: In this case, $\mathbb{E}[n\mathbb{E}[\epsilon_{jt}^{n-1}]\epsilon_{it}^{m}\mathbf{x}_t']=\mathbb{E}[n\mathbb{E}[\epsilon_{jt}^{m}|\mathbf{x}_t]\mathbf{x}_t']=n\mathbb{E}[\epsilon_{jt}^{m}]\mathbb{E}[\mathbf{x}_t']$. Since $n+m\in\{3,4\}$, we either have $m=2$ or $m=3$. In either case, $\mathbb{E}[\epsilon_{jt}^{m}]\neq 0$ in general. The term could thus only be zero if $\mathbb{E}[\mathbf{x}_t]=\mathbf{0}$, which would be satisfied if $\mathbf{y}_t$ were zero mean and no constant term was included. (Alternatively, if we ignored co-skewness, so that $n+m=4$ and considered symmetric distributions, we would get $\mathbb{E}[\epsilon_{jt}^{m}]=0$. That case seems to be too restrictive.)

$n=2$: In this case, $\mathbb{E}[\epsilon_{jt}^{n-1}]=\mathbb{E}[\epsilon_{jt}]=0$ and $\mathbb{E}[n\epsilon_{it}^{m}\epsilon_{jt}^{n-1}\mathbf{x}_t']=\mathbf{0}$.

$n=3$: In this case, $\mathbb{E}[\epsilon_{jt}^{3-1}]=\mathbb{E}[\epsilon_{jt}^{2}]=1$. Now, since $n,m\geq 1$ we must have that $m+n=4$ and $m=1$. Using the law of iterated expectations we get that

$$\mathbb{E}[3\mathbb{E}[\epsilon_{jt}^{3-1}]\epsilon_{it}^{m}\mathbf{x}_t']=\mathbb{E}[3\epsilon_{it}\mathbf{x}_t']=\mathbb{E}[3\mathbb{E}[\epsilon_{it}|\mathbf{x}_t]\mathbf{x}_t']=\mathbb{E}[0\times\mathbf{x}_t']=\mathbf{0}.$$

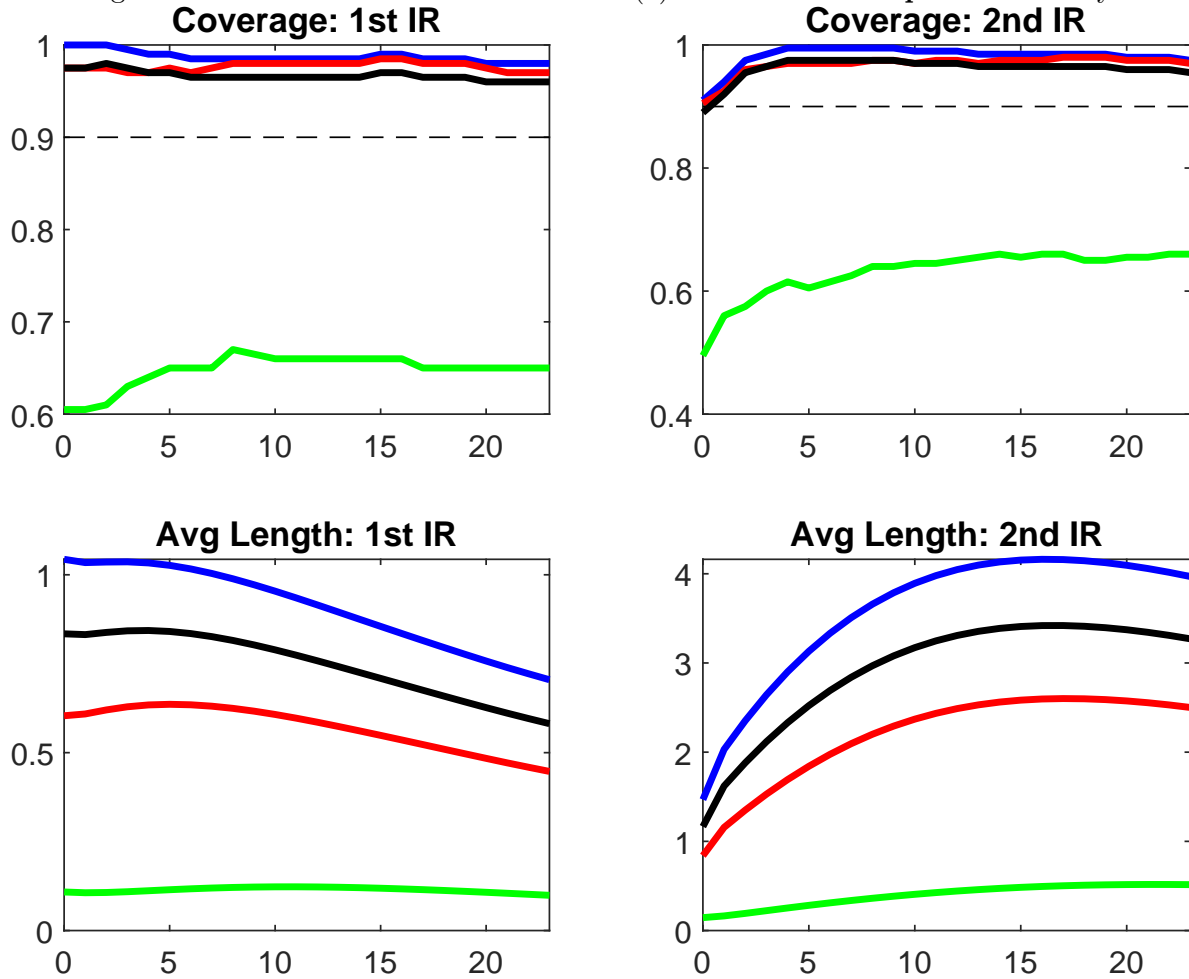If $\mathbb{E}[\mathbf{x}_t] = \mathbf{0}$, we can thus simplify (A.2) as follows:

$$
\frac{1}{\sqrt{T}} \sum_{t=1}^{T} \left(\hat{\epsilon}_{it}^m \hat{\epsilon}_{jt}^n - \epsilon_{it}^m \epsilon_{jt}^n \right) = \left(\frac{1}{T} \sum_{t=1}^{T} m \epsilon_{it}^{m-1} \epsilon_{jt}^n \boldsymbol{\varepsilon}_t' \right) \mathbf{B}^{-1} \underbrace{\sqrt{T}(\hat{\mathbf{B}} - \mathbf{B}_0)}_{\xrightarrow{d} \mathcal{N}} \mathbf{e}_i
$$

$$
+ \left(\frac{1}{T} \sum_{t=1}^{T} n \epsilon_{it}^m \epsilon_{jt}^{n-1} \boldsymbol{\varepsilon}_t' \right) \mathbf{B}^{-1} \underbrace{\sqrt{T}(\hat{\mathbf{B}} - \mathbf{B}_0)}_{\xrightarrow{d} \mathcal{N}} \mathbf{e}_j + O_p(1/\sqrt{T}).
$$

(A.3)

Since in (A.3) the multiplication with $\boldsymbol{\varepsilon}_t' = [\epsilon_{1t}, \ldots, \epsilon_{pt}]$ increases the power of either ϵ_{it} or ϵ_{jt} by one, the RHS does not simplify further.

B   Additional Monte Carlo Simulations

Figure A1: Simulation Results with Skew t(8) Errors and True **q** near Boundary

This figure reports the coverage and average length of impulse responses in the simulated VAR for each horizon in the Monte Carlo simulation with described in the text. Blue Line: Sign Restrictions, Black Line: GMM, Red Line: BKR, Green Line: GMM based on assumption of identification. The sample size is $T = 500$.

Figure A2: Sample size vs size of statistically identified set in Monte Carlo simulations: Standard t_4 distribution vs skewed t_8 distribution.
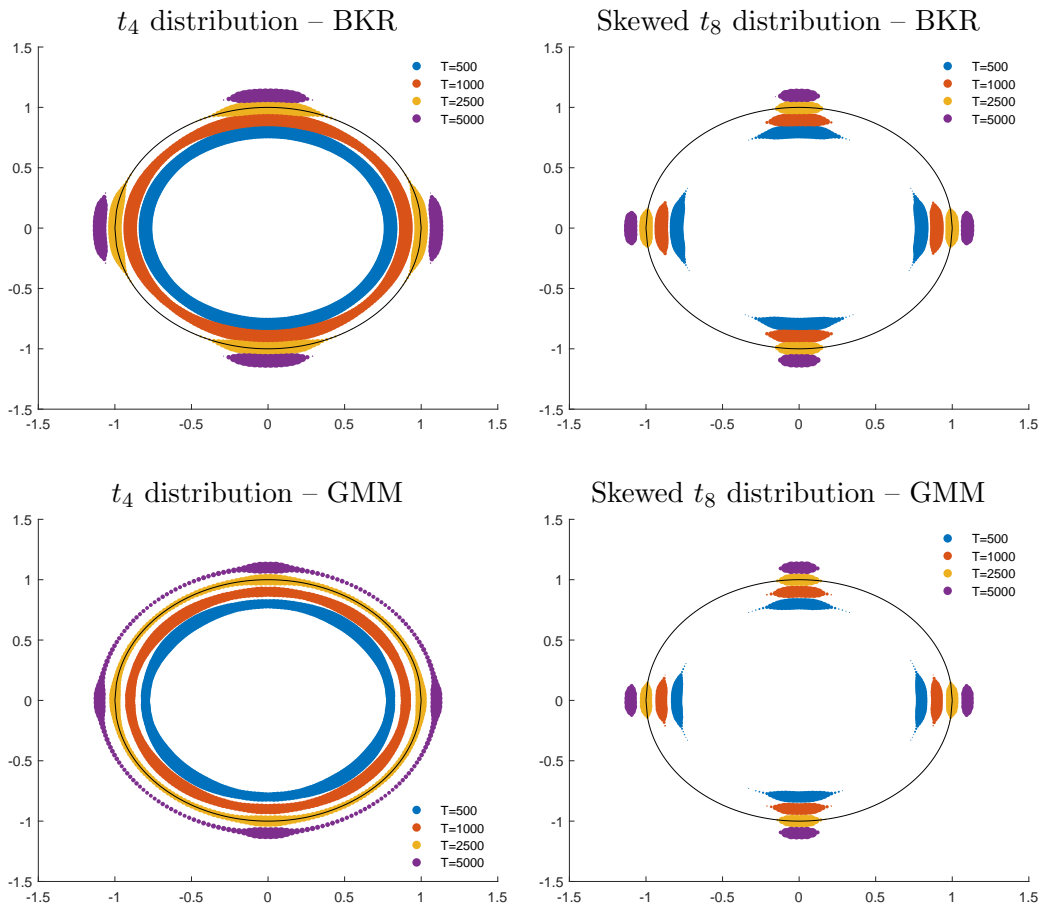
Note: Tests use asymptotic critical values. Width of areas indicates the fraction of simulations in which the segment was included in the statistically identified set.

C Distribution and time series of identified shocks in Bloom (2009)
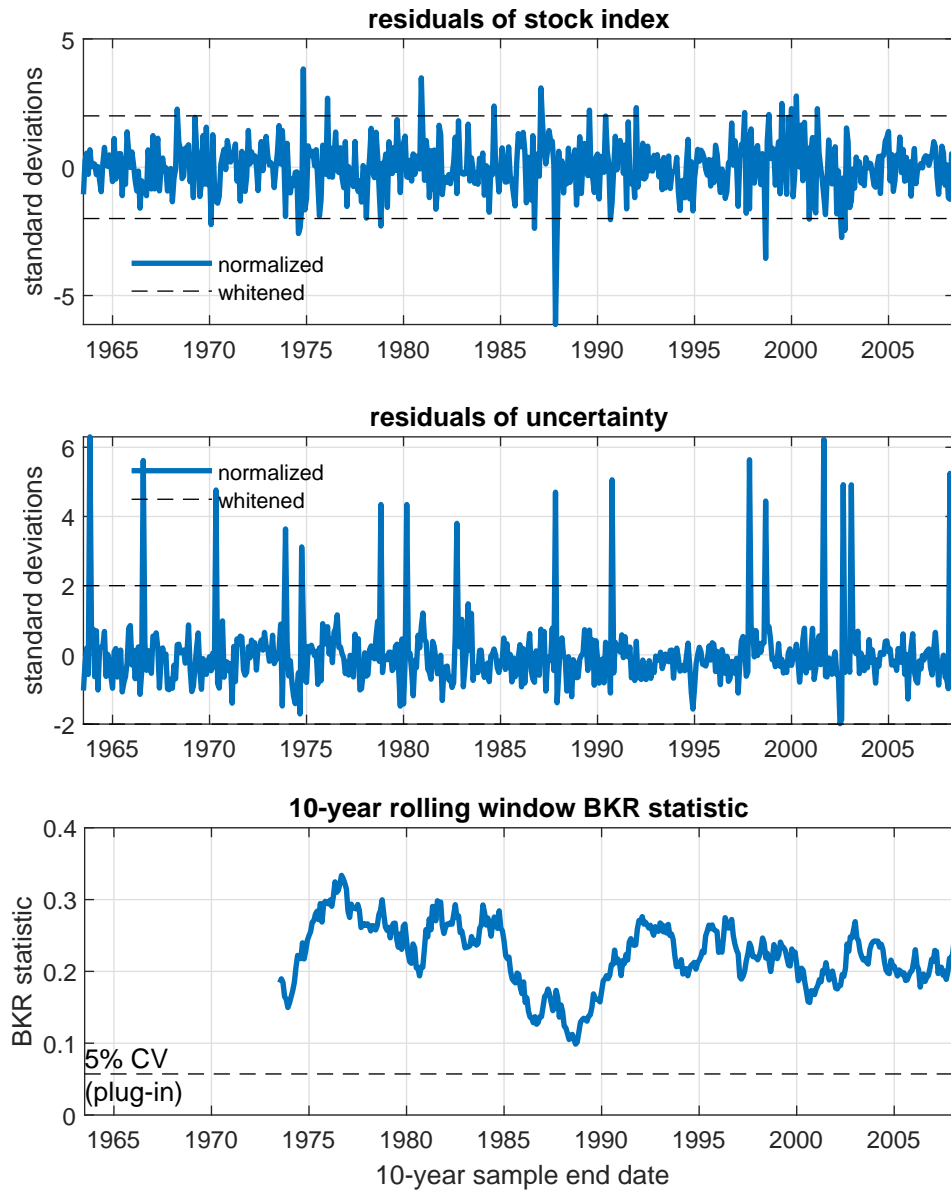
Figure A3: Stock market shocks (top panel), volatility shock (middle panel), and 10-year rolling window of BKR test of identified shocks (bottom panel) in Bloom (2009) application: 1963m6–2008m6
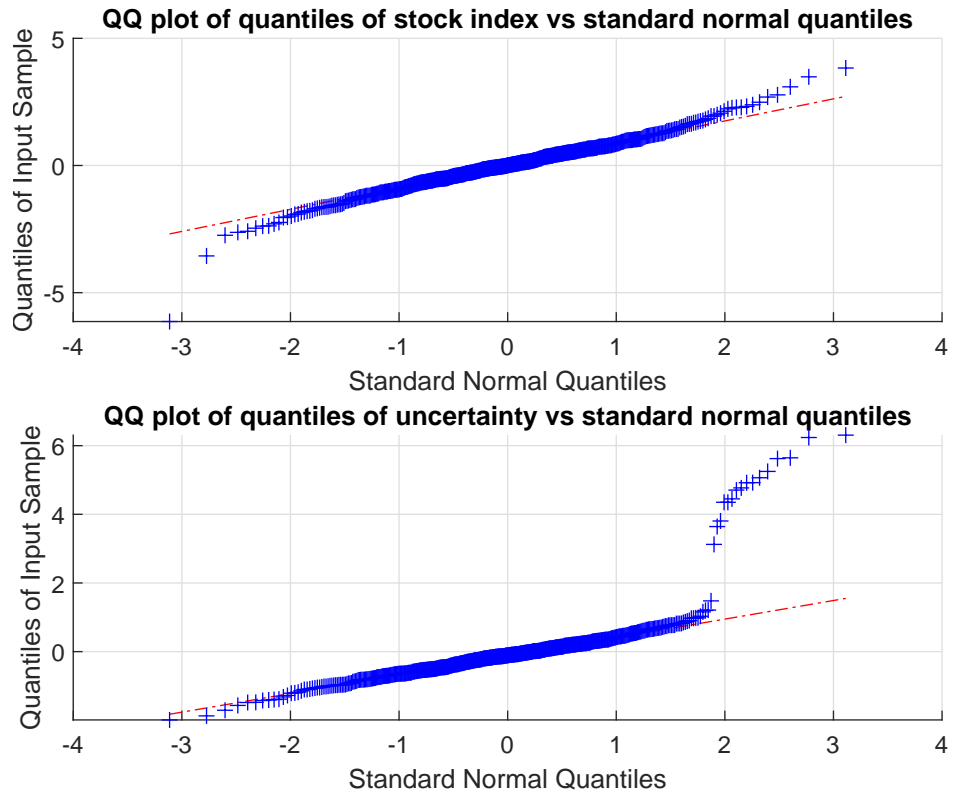
Figure A4: QQ plot of log stock index and uncertainty index shocks in Bloom (2009) application: 1963m6–2008m6

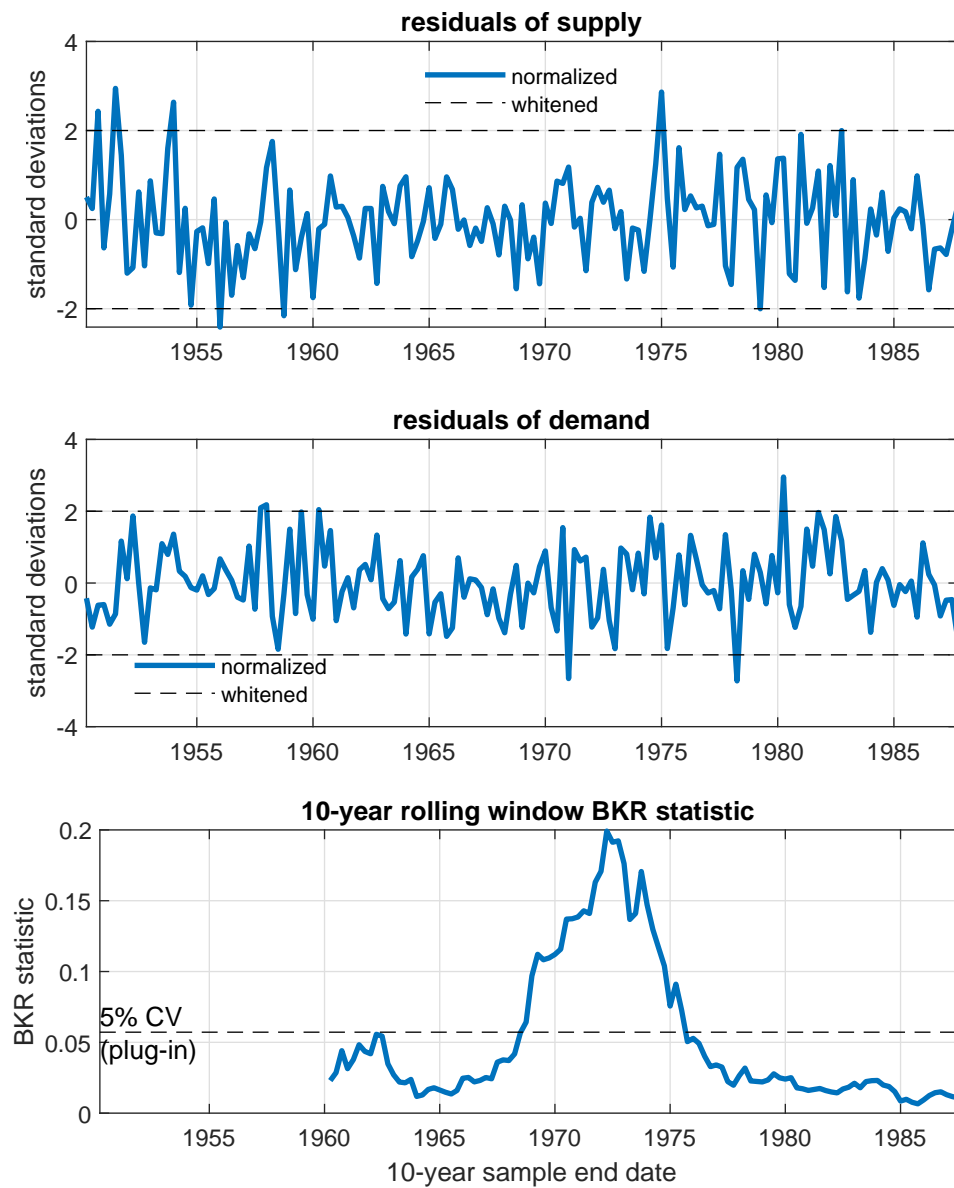# D Distribution and time series of identified shocks in Blanchard and Quah (1989)

Figure A5: Supply shock (top panel), volatility shock (middle panel), and 10-year rolling window of BKR test of identified shocks (bottom panel) in Blanchard and Quah (1989) application: 1950q2–1987q4
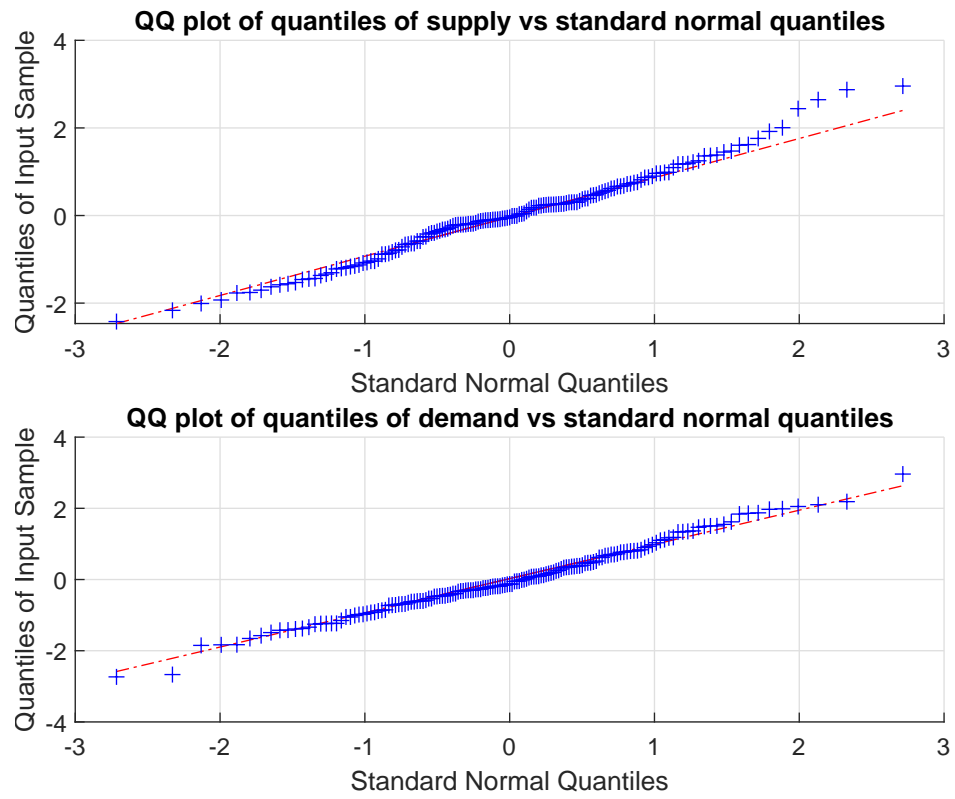
Figure A6: QQ plot of supply and demand shocks in Blanchard and Quah (1989) application: 1950q2–1987q4

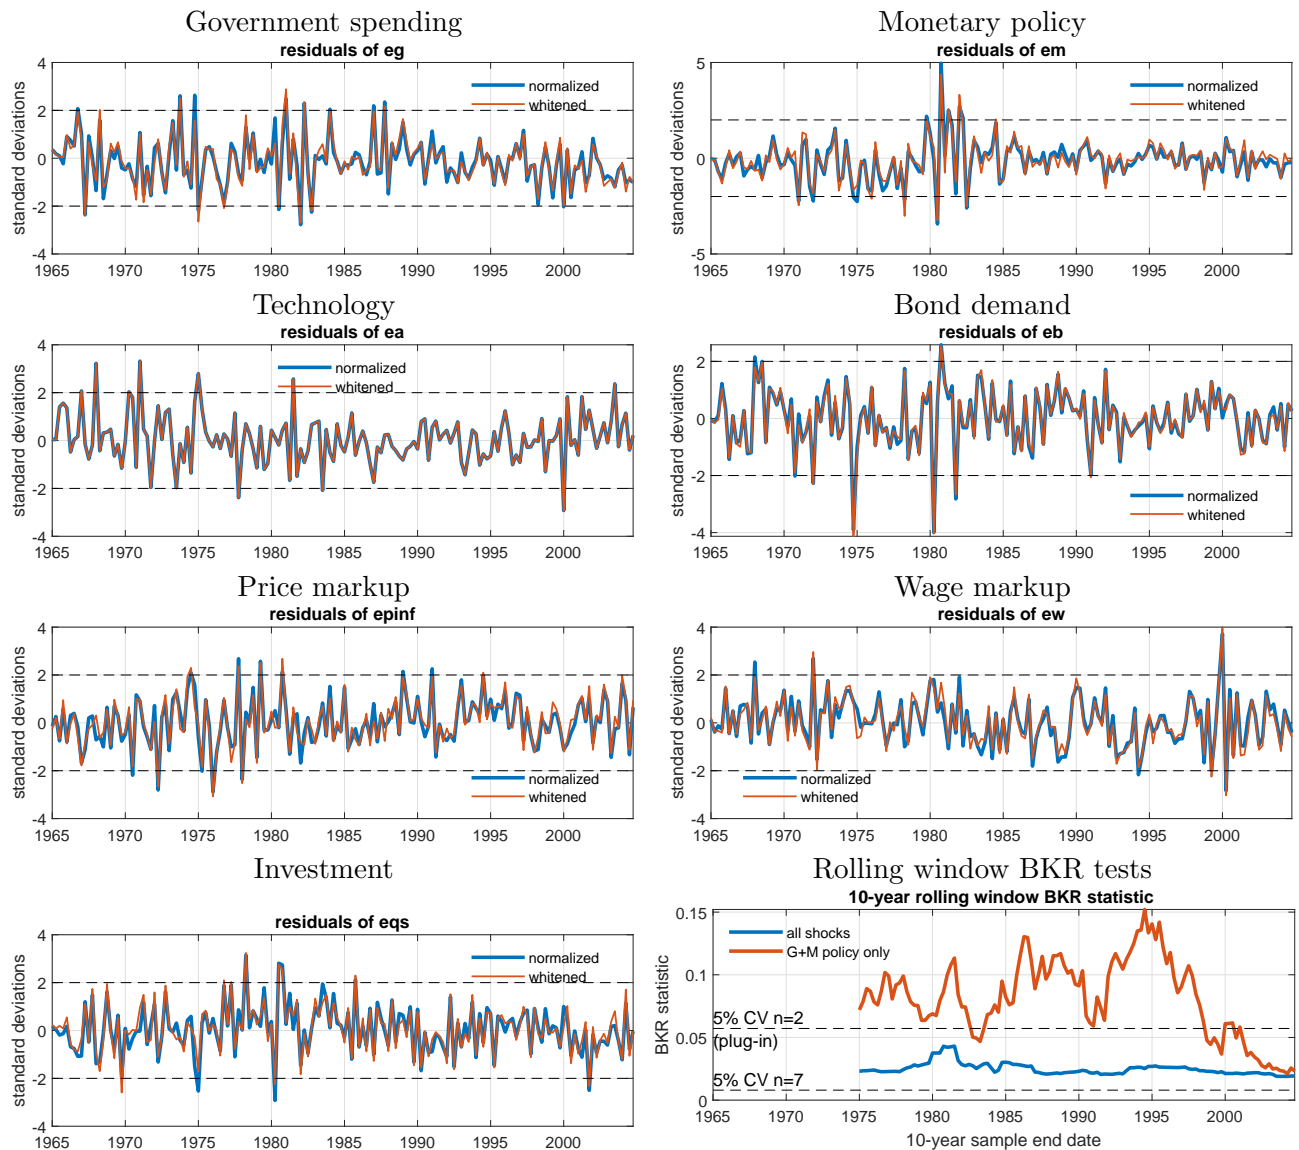# E   Distribution and time series of identified shocks in Smets and Wouters (2007)

Figure A7: Plots of identified shocks in Smets and Wouters (2007) application and 10-year rolling window BKR statistic: 1966q1–2004q4

Government spending

QQ plot of quantiles of eg vs standard normal quantiles

Monetary policy

QQ plot of quantiles of em vs standard normal quantiles

Technology

QQ plot of quantiles of ea vs standard normal quantiles

Bond demand

QQ plot of quantiles of eb vs standard normal quantiles

Price markup

QQ plot of quantiles of epinf vs standard normal quantiles

Wage markup

QQ plot of quantiles of ew vs standard normal quantiles

Investment

QQ plot of quantiles of eqs vs standard normal quantiles

Figure A8: QQ plot of identified shocks in Smets and Wouters (2007) application: 1966q1–2004q4

43

F Data sources and additional information for manufacturing labor market

We use the following definitions:

- Real earnings are defined as average hourly earnings multiplied by the purchasing power of the consumer dollar in U.S. city average: `https://fred.stlouisfed.org/series/CES3000000008` and `https://fred.stlouisfed.org/series/CUUR0000SA0R`.

- Manufacturing employment is all employment in manufacturing: `https://fred.stlouisfed.org/series/MANEMP`

- We transform each series by taking the natural log and multiplying by 100.

For our alternative application to the overall labor market we use the following definitions:

- Real earnings are defined as "Nonfarm Business Sector: Real Hourly Compensation for All Employed Persons": `https://fred.stlouisfed.org/series/COMPRNFB`.

- Employment is "All Employees, Total Nonfarm", aggregated to quarterly data by taking the end of the period value: `https://fred.stlouisfed.org/series/PAYEMS`

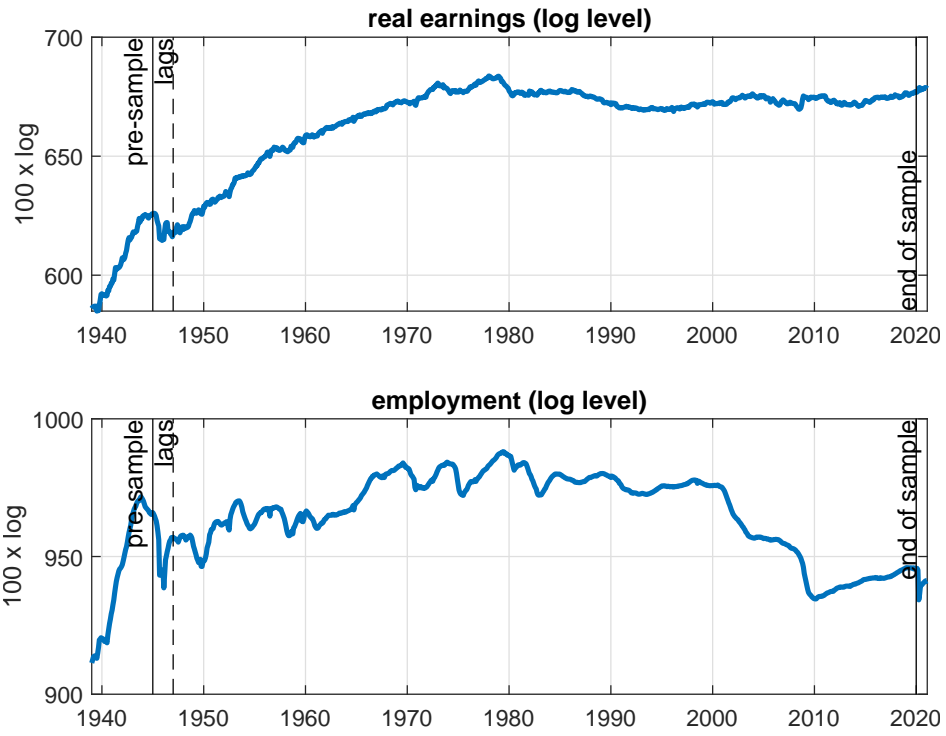- We again transform each series by taking the natural log and multiplying by 100.

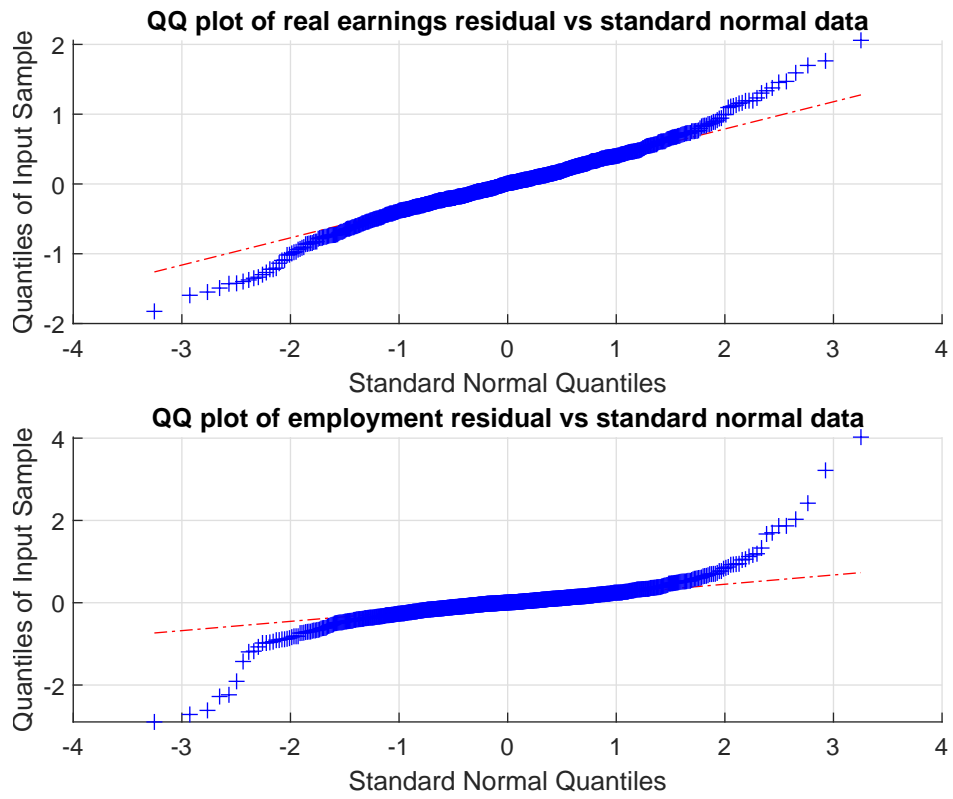Figure A9: Real manufacturing earnings and employment: 1945m1–2019m12

Figure A10: QQ plot of real manufacturing earnings and employment residuals vs standard normal quantiles: 1947m1–2019m12

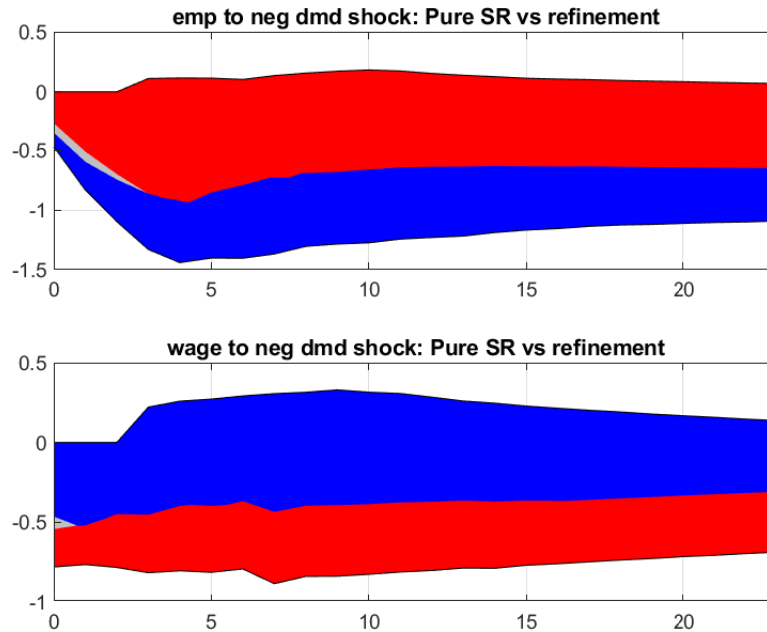## G Refinement of set identification in Baumeister and Hamilton (2015) example

Figure A11: Refinement of set identification in Baumeister and Hamilton (2015) example: Full confidence interval of IRFs

Reduced form shock correlation matrix: $\begin{bmatrix} 1 & \\ 0.06 & 1 \end{bmatrix}$

Figure A12: Refinement of set identification in Baumeister and Hamilton (2015) example: Identified set of Cholesky factor rotations and of impact IRFs at the MLE (Moon and Schorfheide (2012) plot)