FEDERAL RESERVE BANK *of* ATLANTA WORKING PAPER SERIES

Narrative Sign Restrictions for SVARs

Juan Antolín-Díaz and Juan F. Rubio-Ramírez

Working Paper 2016-16a
Revised October 2017

Abstract: We identify structural vector autoregressions using narrative sign restrictions. Narrative sign restrictions constrain the structural shocks and/or the historical decomposition around key historical events, ensuring that they agree with the established narrative account of these episodes. Using models of the oil market and monetary policy, we show that narrative sign restrictions tend to be highly informative. Even a single narrative sign restriction may dramatically sharpen and even change the inference of SVARs originally identified via traditional sign restrictions. Our approach combines the appeal of narrative methods with the popularized usage of traditional sign restrictions.

JEL classification: C32, E52, Q35

Key words: narrative information, SVARs, Bayesian approach, sign restrictions, oil market, monetary policy

---

The authors are grateful to Gavyn Davies, Dan Waggonner, Lutz Kilian, Michele Lenza, Frank Schorfheide, Thomas Drechsel, and Ivan Petrella for helpful comments and suggestions. The views expressed here are the authors' and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. Any remaining errors are the authors' responsibility.

Please address questions regarding content to Juan Antolín-Díaz, Fulcrum Asset Management, 66 Seymour Street, London W1H 5BT, United Kingdom, juan.antolin-diaz@fulcrumasset.com, or Juan F. Rubio-Ramírez (corresponding author), Economics Department, Emory University and Federal Reserve Bank of Atlanta, Rich Memorial Building, Room 306, Atlanta, GA 30322-2240, juan.rubio-ramirez@emory.edu.

1 Introduction

Starting with Faust (1998), Canova and Nicolo (2002), and Uhlig (2005), it has become common to identify structural vector autoregressions (SVARs) using a handful of uncontroversial sign restrictions on either the impulse response functions or the structural parameters themselves. Such minimalist restrictions are generally weaker than classical identification schemes and, therefore, likely to be agreed upon by a majority of researchers. Additionally, because the structural parameters are set-identified, they lead to conclusions that are robust across the set of SVARs that satisfy the sign restrictions (see Rubio-Ramirez et al., 2010 for details). But this minimalist approach is not without cost. The small number of sign restrictions will usually result in a set of structural parameters with very different implications for IRFs, elasticities, historical decompositions or forecasting error variance decompositions. In the best case, this means that it will be difficult to arrive at meaningful economic conclusions. In the worst case, there is the risk of retaining in the admissible set structural parameters with implausible implications. The latter point was first illustrated by Kilian and Murphy (2012), who showed that, in the context of the global market for crude oil, SVARs identified only through sign restrictions on IRFs imply disputable values for the price elasticity of oil supply to demand shocks. More recently, Arias et al. (2016a) have pointed out that the identification scheme of Uhlig (2005) retains many structural parameters with improbable implications for the systematic response of monetary policy to output. The challenge is to come up with a small number of additional uncontentious sign restrictions that help shrink the set of admissible structural parameters and allow us to reach clear economic conclusions.

We propose a new class of sign restrictions based on narrative information that we call narrative sign restrictions. Narrative sign restrictions constrain the structural parameters by ensuring that around selected historical events the structural shocks and/or historical decomposition agree with the established narrative. For example, a narrative sign restriction

on the structural shocks rule out structural parameters that disagree with the view that "a negative oil supply shock occurred at the outbreak of the Gulf War in August 1990", whereas a restriction on the historical decomposition would impose that "a monetary policy shock was the most important driver of the increase in the federal funds rate observed in October 1979." Narrative information in the context of the oil market was used by Kilian and Murphy (2014) to confirm the validity of their proposed identification, but, to the best of our knowledge, we are the first to formalize the idea and develop the methodology. We show that whereas sign restrictions on the IRFs and the structural parameters, which we refer to as traditional sign restrictions, truncate the support of the prior distribution of the structural parameters, narrative sign restrictions instead truncate the support of the likelihood function. Thus, the Bayesian methods in Rubio-Ramirez et al. (2010) and Arias et al. (2016b) need to be modified for the case of narrative sign restrictions. Narrative sign restrictions complement the traditional ones. In our empirical applications we combine both.

A long tradition, starting with Friedman and Schwartz (1963), uses historical sources to identify structural shocks. A key reference is the work of Romer and Romer (1989), who combed through the minutes of the Federal Open Market Committee to single out a number of events that they argued represented monetary policy shocks. A large number of subsequent papers have adopted and extended Romer and Romer's (1989) approach, documenting and collecting various historical events on monetary policy shocks (Romer and Romer, 2004), oil shocks (Hamilton, 1985, Kilian, 2008), and fiscal shocks (Ramey, 2011, Romer and Romer, 2010). The objective of these papers is to construct narrative time series that are then treated as a direct measure of the structural shocks of interest. Recognizing that the narrative time series might be imperfect measures of the structural shocks, recent papers have proposed to treat the narrative time series as external instruments of the targeted structural shocks, i.e., correlated with the shock of interest, and uncorrelated with other structural shocks. This approach was first suggested in Stock and Watson (2008) and was developed independently

by Stock and Watson (2012) and Mertens and Ravn (2013).[1]

There are important differences between our method and the existing narrative approaches. First, in practice our method only uses a small number of key historical events, and sometimes a single event, as opposed to an entire time series. Like the instrumental variables approach, this alleviates the issue of measurement error in the narrative time series, but with our method the researcher can incorporate only those events upon which there is agreement. It also makes it straightforward to verify how a particular episode affects the results. Second, we impose the narrative information as sign restrictions. For instance, one might not be sure of exactly how much of the October 1979 Volcker reform was exogenous, but one is confident that a contractionary monetary policy shock did occur, and that it was more relevant than other shocks in explaining the unexpected movement in the federal funds rate. Therefore, our method combines the appeal of narrative approaches with the advantages of sign restrictions. Finally, our methods are Bayesian, while most of the existing narrative approaches are frequentist.

We illustrate the methodology by applying it to two well-known examples of SVARs previously identified with traditional sign restrictions for which narrative information is readily available. In particular, we revisit the model of the oil market of Kilian and Murphy (2012) and Inoue and Kilian (2013), and the model of the effects of monetary policy that has been used in Christiano et al. (1999), Bernanke and Mihov (1998), and Uhlig (2005). In the case of oil shocks, supply shocks are sharply identified using only traditional sign restrictions, whereas disentangling aggregate demand and oil-specific demand shocks is more difficult in a standard three-variable oil market VAR. Adding narrative sign restrictions based on a small set of historical events dramatically sharpens the identification. In fact, adding narrative information on a single event, the start of the Persian Gulf War in August 1990, is enough to obtain this result. In the case of monetary policy shocks, we show that Uhlig's (2005)

[1] See also Montiel-Olea et al. (2015).

results are not robust to discarding structural parameters that have implausible implications for the key historical event that occurred in October of 1979, the Volcker reform. In both applications, we find that restrictions on the historical decomposition tend to be particularly effective in shrinking the identified set.

The rest of this paper is organized as follows. Section 2 presents the basic SVAR framework. Section 3 introduces narrative sign restrictions. Section 4 derives the posterior distribution under narrative sign restrictions and describes the algorithm to draw from it. Sections 5 and 6 apply the methodology to the oil market and of monetary policy shocks respectively. Section 7 concludes.

2 The Model

Consider the structural vector autoregression (SVAR) of the general form

$$\mathbf{y}_t'\mathbf{A}_0 = \sum_{\ell=1}^{p} \mathbf{y}_{t-\ell}'\mathbf{A}_\ell + \mathbf{c} + \varepsilon_t' \quad \text{for } 1 \leq t \leq T \tag{1}$$

where $\mathbf{y}_t$ is an $n \times 1$ vector of variables, ε_t is an $n \times 1$ vector of structural shocks, $\mathbf{A}_\ell$ is an $n \times n$ matrix of parameters for $0 \leq \ell \leq p$ with $\mathbf{A}_0$ invertible, $\mathbf{c}$ is a $1 \times n$ vector of parameters, p is the lag length, and T is the sample size. The vector ε_t, conditional on past information and the initial conditions $\mathbf{y}_0, \ldots, \mathbf{y}_{1-p}$, is Gaussian with mean zero and covariance matrix $\mathbf{I}_n$, the $n \times n$ identity matrix. The model described in Equation (1) can be written as

$$\mathbf{y}_t'\mathbf{A}_0 = \mathbf{x}_t'\mathbf{A}_+ + \varepsilon_t' \quad \text{for } 1 \leq t \leq T, \tag{2}$$

where $\mathbf{A}_+' = \begin{bmatrix}\mathbf{A}_1' & \cdots & \mathbf{A}_p' & \mathbf{c}'\end{bmatrix}$ and $\mathbf{x}_t' = \begin{bmatrix}\mathbf{y}_{t-1}', \ldots, \mathbf{y}_{t-p}', 1\end{bmatrix}$ for $1 \leq t \leq T$. The dimension of $\mathbf{A}_+$ is $m \times n$ and the dimension of $\mathbf{x}_t$ is $m \times 1$, where $m = np + 1$. The reduced-form representation implied by Equation (2) is $\mathbf{y}_t' = \mathbf{x}_t'\mathbf{B} + \mathbf{u}_t'$ for $1 \leq t \leq T$, where $\mathbf{B} = \mathbf{A}_+\mathbf{A}_0^{-1}$,

$\mathbf{u}_t' = \varepsilon_t' \mathbf{A}_0^{-1}$, and $\mathbb{E}\left[\mathbf{u}_t \mathbf{u}_t'\right] = \Sigma = \left(\mathbf{A}_0 \mathbf{A}_0'\right)^{-1}$. The matrices $\mathbf{B}$ and Σ are the reduced-form parameters, while $\mathbf{A}_0$ and $\mathbf{A}_+$ are the structural parameters. Similarly, $\mathbf{u}_t'$ are the reduced-form innovations, while ε_t' are the structural shocks. The shocks are orthogonal and have an economic interpretation, while the innovations are, in general, correlated and do not have an interpretation. Let $\mathbf{\Theta} = (\mathbf{A}_0, \mathbf{A}_+)$ collect the value of the structural parameters.

2.1 Impulse response functions

Recall the definition of impulse response functions (IRFs). Given a value $\mathbf{\Theta}$ of the structural parameters, the response of the i-th variable to the j-th structural shock at horizon k corresponds to the element in row i and column j of the matrix $\mathbf{L}_k(\mathbf{\Theta})$, where $\mathbf{L}_k(\mathbf{\Theta})$ is defined recursively by

$$\mathbf{L}_0(\mathbf{\Theta}) = \left(\mathbf{A}_0^{-1}\right)', \quad \mathbf{L}_k(\mathbf{\Theta}) = \sum_{\ell=1}^{k} \left(\mathbf{A}_\ell \mathbf{A}_0^{-1}\right)' \mathbf{L}_{k-\ell}(\mathbf{\Theta}), \text{ for } 1 \leq k \leq p,$$

$$\mathbf{L}_k(\mathbf{\Theta}) = \sum_{\ell=1}^{p} \left(\mathbf{A}_\ell \mathbf{A}_0^{-1}\right)' \mathbf{L}_{k-\ell}(\mathbf{\Theta}), \text{ for } p < k < \infty.$$

2.2 Structural shocks and historical decomposition

Given a value $\mathbf{\Theta}$ of the structural parameters and the data, the structural shocks at time t are

$$\varepsilon_t'(\mathbf{\Theta}) = \mathbf{y}_t' \mathbf{A}_0 - \mathbf{x}_t' \mathbf{A}_+ \text{ for } 1 \leq t \leq T. \tag{3}$$

The historical decomposition calculates the cumulative contribution of each shock to the observed unexpected change in the variables between two periods.[2] Formally, the contribution of the j-th shock to the observed unexpected change in the i-th variable between periods t

[2]See Kilian and Lütkepohl (2017) for a textbook treatment.

and $t + h$ is

$$H_{i,j,t,t+h}(\mathbf{\Theta}, \varepsilon_t, \ldots, \varepsilon_{t+h}) = \sum_{\ell=0}^{h} \mathbf{e}'_{i,n}\mathbf{L}_\ell(\mathbf{\Theta})\mathbf{e}_{j,n}\mathbf{e}'_{j,n}\varepsilon_{t+h-\ell},$$

where $\mathbf{e}_{j,n}$ is the j-th column of $\mathbf{I}_n$, for $1 \leq i, j \leq n$ and for $h \geq 0$.

3 The Identification Problem and Sign Restrictions

As is well known, the structural form in Equation (1) is not identified, so restrictions must be imposed on the structural parameters to solve the identification problem. The desire to impose only minimalist identification restrictions that are agreed upon by most researchers and lead to robust conclusions motivated Faust (1998), Canova and Nicolo (2002) and Uhlig (2005) to develop methods to identify the structural parameters by placing a handful of uncontroversial sign restrictions on the IRFs or the structural parameters themselves. In this paper we propose a new class of sign restrictions based on narrative information that we call narrative sign restrictions. Narrative sign restrictions constrain the structural parameters by ensuring that around a handful of key historical events the structural shocks and/or historical decompositions agree with the established narrative. For instance, in the context of a model of demand and supply in the global oil market, it is well established from historical sources that an exogenous disruption to oil production occurred at the outbreak of the Gulf War in August 1990. Therefore a researcher may want to constrain the structural parameters so that the oil supply shock for that period was negative or that it was the most important contributor (as opposed to, for instance, a negative demand shock) to the unexpected drop in oil production observed during that period. We now formally describe the functions that characterize sign restrictions on the IRFs and the structural parameters (traditional sign restrictions) and on the structural shocks and the historical decompositions (narrative sign restrictions).

3.1 Traditional sign restrictions

Traditional sign restrictions are well understood and their use is widespread in the literature. In particular, Rubio-Ramirez et al. (2010) and Arias et al. (2016b) highlight how this class of restrictions can be characterized by the function

$$\Gamma\left(\mathbf{\Theta}\right) = \left(\mathbf{e}_{1,n}'\mathbf{F}(\mathbf{\Theta})'\mathbf{S}_1', \cdots, \mathbf{e}_{n,n}'\mathbf{F}(\mathbf{\Theta})'\mathbf{S}_n'\right)' > \mathbf{0}. \tag{4}$$

Appropriate choices of $\mathbf{S}_j$ and $\mathbf{F}(\mathbf{\Theta})$ will lead to sign restrictions on the IRFs or the structural parameters themselves. In particular, to impose restrictions on the IRFs, one can define $\mathbf{F}(\mathbf{\Theta})$ as vertically stacking the IRFs at the different horizons over which we want to impose the restrictions and $\mathbf{S}_j$ as an $s_j \times r_j$ matrix of zeros, ones and negative ones that will select the horizons and the variables over which we want to impose the r_j sign restrictions to identify structural shock j. If instead we want to impose restrictions on the structural parameters themselves, we can then define $\mathbf{F}(\mathbf{\Theta}) = \mathbf{\Theta}$ and $\mathbf{S}_j$ as an $s_j \times r_j$ matrix of zeros, ones and negative ones that will select entries of $\mathbf{\Theta}$ over which we want to impose the sign restrictions.

3.2 Restrictions on the signs of the structural shocks

Let us now consider the first class of narrative sign restrictions. Let us assume that we want to impose the restriction that the signs of the j-th shock at s_j episodes occurring at dates $t_1, \ldots, t_{s_j}$ are all positive. Then, the narrative sign restrictions can be imposed as

$$\mathbf{e}_{j,n}'\varepsilon_{t_v}\left(\mathbf{\Theta}\right) > 0 \text{ for } 1 \leq v \leq s_j. \tag{5}$$

Assume instead that we want to impose the restriction that the signs of the j-th shock at s_j episodes occurring at dates $t_1, \ldots, t_{s_j}$ are negative. Then, the narrative sign restrictions can be imposed with a negative sign in the left-hand side of Equation (5). Of course, one could

restrict the shocks in a few periods to be negative and positive in a few others.

3.3 Restrictions on the historical decomposition

Let us now consider the second class of narrative sign restrictions. In many cases the researcher will have narrative information that indicates that a particular shock was the most important contributor to the unexpected movement of some variable during a particular period. This is information on the relative magnitude of the contribution of the j-th shock to the unexpected change in the i-th variable between some periods, i.e. on the historical decomposition. We propose to formalize this idea in two different ways. First, we may specify that a given shock was the *most important* (*least important*) driver of the unexpected change in a variable during some periods. By this we mean that for a particular period or periods the absolute value of its contribution to the unexpected change in a variable is larger (smaller) than the absolute value of the contribution of any other structural shock. Second, we may want to say that a given shock was the *overwhelming* (*negligible*) driver of the unexpected change in a given variable during the period. By this we mean that for a particular period or periods the absolute value of its contribution to the unexpected change in a variable is larger (smaller) than the sum of the absolute value of the contributions of all other structural shocks. We will label these two alternatives Type A and Type B, respectively.[3]

3.4 Type A restrictions on the historical decomposition

To fix ideas, consider the following example: assume we have a model with three variables and we want to impose that between periods 6 and 7, the 2nd structural shock is the

[3]As pointed to us by a referee, one could also impose sign restrictions on the historical decompositions themselves, rather than on their relative magnitudes. For example, Kilian and Murphy (2014) note that industry sources show that the cumulative effect of speculative demand shocks between May 1979 and December 1979 on the real price of oil was positive, without this effect necessarily being the dominant effect. This type of restriction would be weaker than any of the three proposed above.

most important contributor in absolute terms to the unexpected change in the 3rd variable. This narrative restriction can be formalized by the function $|H_{3,2,6,7}(\mathbf{\Theta}, \varepsilon_6(\mathbf{\Theta}), \varepsilon_7(\mathbf{\Theta}))| - \max_{j' \neq 2} |H_{3,j',6,7}(\mathbf{\Theta}, \varepsilon_6(\mathbf{\Theta}), \varepsilon_7(\mathbf{\Theta}))| > 0$, where $|H(\cdot)|$ is the absolute value of the matrix $H(\cdot)$. In other words, the contribution of the 2nd shock to the historical decomposition is larger in absolute value than the largest contribution of any other shock.

In general, we can identify the j-th shock by imposing s_j restrictions of this type. Thus, suppose we want to impose the restriction that the j-th shock is the *most important* contributor to the unexpected change in the $i_1, \ldots, i_{s_j}$-th variables from periods $t_1, \ldots, t_{s_j}$ to $t_1 + h_1, \ldots, t_{s_j} + h_{s_j}$, i.e. that its cumulative contribution is larger in absolute value than the contribution of any other shock to the unexpected change in those variables during those periods. Then, the narrative sign restrictions can be imposed as

$$|H_{i_v,j,t_v,t_v+h_v}(\mathbf{\Theta}, \varepsilon_{t_v}(\mathbf{\Theta}), \ldots, \varepsilon_{t_v+h_v}(\mathbf{\Theta}))| - \max_{j' \neq j} |H_{i_v,j',t_v,t_v+h_v}(\mathbf{\Theta}, \varepsilon_{t_v}(\mathbf{\Theta}), \ldots, \varepsilon_{t_v+h_v}(\mathbf{\Theta}))| > 0, \quad (6)$$

for $1 \leq v \leq s_j$. If instead one wishes to impose that the contribution of the shocks is the *least important*, i.e. that its cumulative contribution is smaller in absolute value than the contribution of any other shock to the unexpected change in those variables during those periods, the narrative sign restrictions can be imposed as

$$|H_{i_v,j,t_v,t_v+h_v}(\mathbf{\Theta}, \varepsilon_{t_v}(\mathbf{\Theta}), \ldots, \varepsilon_{t_v+h_v}(\mathbf{\Theta}))| - \min_{j' \neq j} |H_{i_v,j',t_v,t_v+h_v}(\mathbf{\Theta}, \varepsilon_{t_v}(\mathbf{\Theta}), \ldots, \varepsilon_{t_v+h_v}(\mathbf{\Theta}))| < 0, \quad (7)$$

for $1 \leq v \leq s_j$. As above, Equations (6) and (7) can be used jointly.

3.4.1 Type B restrictions on the historical decomposition

As before, to fix ideas, assume we have a model with three variables and we want to impose the restriction that between periods 6 and 7, the 2nd structural shock is the overwhelming contribu-

tor in absolute terms to the unexpected change in the 3rd variable. This narrative restriction can be formalized by the function $|H_{3,2,6,7}(\mathbf{\Theta}, \varepsilon_6(\mathbf{\Theta}), \varepsilon_7(\mathbf{\Theta}))| - \sum_{j' \neq 2} |H_{3,j',6,7}(\mathbf{\Theta}, \varepsilon_6(\mathbf{\Theta}), \varepsilon_7(\mathbf{\Theta}))| > 0$. In other words, the contribution of the 2nd shock to the historical decomposition is larger in absolute value than the sum of the absolute contributions of all other shocks.

As before, we can identify the j-th structural shock by imposing s_j restrictions of this type. Thus, suppose we want to impose the restriction that the j-th shock is the *overwhelming* contributor to the unexpected change in the $i_1, \ldots, i_{s_j}$-th variables from periods $t_1, \ldots, t_{s_j}$ to $t_1 + h_1, \ldots, t_{s_j} + h_{s_j}$, i.e. that its contribution is larger in absolute value than the sum of the absolute contributions of all other shocks to the unexpected change in those variables during those periods. Then, we can define

$$|H_{i_v,j,t_v,t_v+h_v}(\mathbf{\Theta}, \varepsilon_{t_v}(\mathbf{\Theta}), \ldots, \varepsilon_{t_v+h_v}(\mathbf{\Theta}))| - \sum_{j' \neq j} |H_{i_v,j',t_v,t_v+h_v}(\mathbf{\Theta}, \varepsilon_{t_v}(\mathbf{\Theta}), \ldots, \varepsilon_{t_v+h_v}(\mathbf{\Theta}))| > 0, \quad (8)$$

for $1 \leq v \leq s_j$. If instead one wishes to impose that the contribution of the shocks is *negligible*, i.e. that its contribution is smaller in absolute value than the sum of the contributions of all other shocks to the unexpected change in those variables during those periods, the narrative sign restrictions can be imposed as

$$|H_{i_v,j,t_v,t_v+h_v}(\mathbf{\Theta}, \varepsilon_{t_v}(\mathbf{\Theta}), \ldots, \varepsilon_{t_v+h_v}(\mathbf{\Theta}))| - \sum_{j' \neq j} |H_{i_v,j',t_v,t_v+h_v}(\mathbf{\Theta}, \varepsilon_{t_v}(\mathbf{\Theta}), \ldots, \varepsilon_{t_v+h_v}(\mathbf{\Theta}))| < 0, \quad (9)$$

for $1 \leq v \leq s_j$. Equations (8) and (9) can also be used jointly.

3.4.2 Discussion

A natural question is to ask whether Type A or Type B restrictions on the historical decomposition are more restrictive. The answer depends on whether we are restricting the cumulative contribution of a particular shock to the unexpected change in a variable to

be "larger" or "smaller." If the contribution of shock j is larger than the sum of all other contributions, it is always larger than any single contribution. Therefore, when contributions are defined as "larger," Type B is more restrictive than Type A. In contrast, if the contribution of shock j is smaller than any single contribution, it must also be smaller than the sum of the other contributions in absolute value. Consequently, when restrictions are defined as "smaller," Type B is stronger than Type A. Therefore, the use of either Type A or Type B allows the researcher to express different levels of confidence in the narrative information about a particular episode.

4 Bayesian Inference

In this section we show how to adapt the Bayesian methods developed in Rubio-Ramirez et al. (2010) and Arias et al. (2016b) to handle narrative sign restrictions. Equations (5)-(9) imply the following function to characterize narrative sign restrictions

$$\phi\left(\mathbf{\Theta}, \varepsilon^{v}\right) > \mathbf{0}, \tag{10}$$

where $\varepsilon^{v} = \left(\varepsilon_{t_1}, \ldots, \varepsilon_{t_v}\right)$ are the structural shocks constrained by the narrative sign restrictions. A comparison with Equation (4) makes it clear that the traditional sign restrictions depend on the structural parameters, whereas the narrative sign restrictions depend as well on the structural shocks. Moreover, Equation (3) implies the following invertible function

$$\varepsilon_t = g_h\left(\mathbf{y}_t, \mathbf{x}_t, \mathbf{\Theta}\right) \text{ for } 1 \leq t \leq T, \tag{11}$$

with $\mathbf{y}_t = g_h^{-1}\left(\varepsilon_t; \mathbf{x}_t, \mathbf{\Theta}\right)$ for $1 \leq t \leq T$. Using Equations (10) and (11), we can write

$$\tilde{\phi}\left(\mathbf{\Theta}, \mathbf{y}^{v}, \mathbf{x}^{v}\right) = \phi\left(\mathbf{\Theta}, g_h\left(\mathbf{y}_{t_1}, \mathbf{x}_{t_1}, \mathbf{\Theta}\right), \ldots, g_h\left(\mathbf{y}_{t_v}, \mathbf{x}_{t_v}, \mathbf{\Theta}\right)\right) > \mathbf{0}, \tag{12}$$

where $\mathbf{y}^v = \left(\mathbf{y}_{t_1}, \ldots, \mathbf{y}_{t_v}\right)$ and $\mathbf{x}^v = \left(\mathbf{x}_{t_1}, \ldots, \mathbf{x}_{t_v}\right)$. Hence, given the data, Equation (10) is continuous on the structural parameters while, given the structural parameters, Equation (10) is continuous on the structural shocks.

4.1 The posterior distribution

Following Arias et al. (2016b), we can consider an alternative parameterization of the structural VAR in (2), defined by $\mathbf{B}$, Σ, and $\mathbf{Q}$, where $\mathbf{Q} \in O(n)$, the set of all orthogonal $n \times n$ matrices, which we call the orthogonal reduced-form parameterization. To define a mapping between $\mathbf{\Theta}$ and $(\mathbf{B}, \Sigma, \mathbf{Q})$, one must first choose a decomposition of the covariance matrix Σ. Let $h(\Sigma)$ be an $n \times n$ matrix that satisfies $h(\Sigma)'h(\Sigma) = \Sigma$, where h is differentiable. One would normally choose $h(\Sigma)$ to be the Cholesky decomposition. Given a decomposition h, we can define the mapping between $\mathbf{\Theta}$ and $(\mathbf{B}, \Sigma, \mathbf{Q})$

$$f_h(\mathbf{\Theta}) = (\underbrace{\mathbf{A}_+\mathbf{A}_0^{-1}}_{\mathbf{B}}, \underbrace{(\mathbf{A}_0\mathbf{A}_0')^{-1}}_{\Sigma}, \underbrace{h((\mathbf{A}_0\mathbf{A}_0')^{-1})\mathbf{A}_0}_{\mathbf{Q}}),$$

where it is easy to see that $h((\mathbf{A}_0\mathbf{A}_0')^{-1})\mathbf{A}_0$ is an orthogonal matrix. The function f_h is invertible, with inverse defined by

$$f_h^{-1}(\mathbf{B}, \Sigma, \mathbf{Q}) = (\underbrace{h(\Sigma)^{-1}\mathbf{Q}}_{\mathbf{A}_0}, \underbrace{\mathbf{B}h(\Sigma)^{-1}\mathbf{Q}}_{\mathbf{A}_+}). \tag{13}$$

Using Equation (13), we can rewrite Equation (12) as $\Phi(\mathbf{B}, \Sigma, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) = \tilde{\phi}\left(f_h^{-1}(\mathbf{B}, \Sigma, \mathbf{Q}), \mathbf{y}^v, \mathbf{x}^v\right) > \mathbf{0}$.

Thus, the posterior of $(\mathbf{B}, \Sigma, \mathbf{Q})$ subject to the narrative sign restrictions is

$$\pi\left(\mathbf{B}, \Sigma, \mathbf{Q}|\mathbf{y}^T, \Phi(\mathbf{B}, \Sigma, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}\right) = \frac{\pi\left(\mathbf{y}^T|\mathbf{B}, \Sigma, \mathbf{Q}, \Phi(\mathbf{B}, \Sigma, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}\right)\pi\left(\mathbf{B}, \Sigma, \mathbf{Q}\right)}{\int \pi\left(\mathbf{y}^T|\mathbf{B}, \Sigma, \mathbf{Q}, \Phi(\mathbf{B}, \Sigma, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}\right)\pi\left(\mathbf{B}, \Sigma, \mathbf{Q}\right) d\left(\mathbf{B}, \Sigma, \mathbf{Q}\right)}, \tag{14}$$

where $\mathbf{y}^T = \{\mathbf{y}_{1-p}, \ldots, \mathbf{y}_0, \ldots, \mathbf{y}_T\}$ is the data, $\pi\left(\mathbf{y}^T|\mathbf{B}, \Sigma, \mathbf{Q}, \Phi(\mathbf{B}, \Sigma, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}\right)$ is the likelihood function subject to the narrative sign restrictions and $\pi\left(\mathbf{B}, \Sigma, \mathbf{Q}\right)$ is the prior.

It is useful at this point to compare the posterior distribution defined in Equation (14) with the one obtained using only traditional sign restrictions. The posterior of $(\mathbf{B}, \Sigma, \mathbf{Q})$ subject to the traditional sign restrictions is

$$\pi\left(\mathbf{B}, \Sigma, \mathbf{Q}|\mathbf{y}^T, \Gamma\left(f_h^{-1}(\mathbf{B}, \Sigma, \mathbf{Q})\right) > \mathbf{0}\right) = \frac{\pi\left(\mathbf{y}^T|\mathbf{B}, \Sigma\right)\pi\left(\mathbf{B}, \Sigma, \mathbf{Q}|\Gamma\left(f_h^{-1}(\mathbf{B}, \Sigma, \mathbf{Q})\right) > \mathbf{0}\right)}{\int \pi\left(\mathbf{y}^T|\mathbf{B}, \Sigma\right)\pi\left(\mathbf{B}, \Sigma, \mathbf{Q}|\Gamma\left(f_h^{-1}(\mathbf{B}, \Sigma, \mathbf{Q})\right) > \mathbf{0}\right) d(\mathbf{B}, \Sigma, \mathbf{Q})},$$

where $\pi\left(\mathbf{y}^T|\mathbf{B}, \Sigma\right)$ is the likelihood function and $\pi\left(\mathbf{B}, \Sigma, \mathbf{Q}|\Gamma\left(f_h^{-1}(\mathbf{B}, \Sigma, \mathbf{Q})\right) > \mathbf{0}\right)$ is the prior subject to the traditional sign restrictions. Since the likelihood function does not depend on $\mathbf{Q}$ and the traditional sign restrictions are characterized by a function that does not depend on the structural shocks, traditional sign restrictions only truncate the prior of $(\mathbf{B}, \Sigma, \mathbf{Q})$. On the contrary, since the function characterizing the narrative sign restrictions depends on the structural shocks, narrative sign restrictions do not truncate the prior of $(\mathbf{B}, \Sigma, \mathbf{Q})$ but the likelihood function.

The truncated likelihood function in Equation (14) can be written as

$$\pi\left(\mathbf{y}^T|\mathbf{B}, \Sigma, \mathbf{Q}, \Phi(\mathbf{B}, \Sigma, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}\right) = \frac{\left[\Phi(\mathbf{B}, \Sigma, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}\right]\pi\left(\mathbf{y}^T|\mathbf{B}, \Sigma\right)}{\int \left[\Phi(\mathbf{B}, \Sigma, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}\right]\pi\left(\mathbf{y}^T|\mathbf{B}, \Sigma\right) d\mathbf{y}^T}. \tag{15}$$

But note that

$$\int \left[\Phi(\mathbf{B},\Sigma,\mathbf{Q},\mathbf{y}^v,\mathbf{x}^v) > \mathbf{0}\right] \pi\left(\mathbf{y}^T|\mathbf{B},\Sigma\right) d\mathbf{y}^T = \int \left[\Phi(\mathbf{B},\Sigma,\mathbf{Q},\mathbf{y}^v,\mathbf{x}^v) > \mathbf{0}\right] \left(\prod_{t=1}^{T} \pi\left(\mathbf{y}_t|\mathbf{x}_t,\mathbf{B},\Sigma\right)\right) d\left(\mathbf{y}_1 \ldots \mathbf{y}_T\right)$$

$$= \int \left[\tilde{\Phi}\left(\mathbf{B},\Sigma,\mathbf{Q},\varepsilon^v\right) > \mathbf{0}\right] \left(\prod_{t=1}^{T} \frac{\pi\left(g_h^{-1}\left(\varepsilon_t;\mathbf{x}_t,f_h^{-1}\left(\mathbf{B},\Sigma,\mathbf{Q}\right)\right)|\mathbf{x}_t,\mathbf{B},\Sigma\right)}{\upsilon_{g_h}\left(g_h^{-1}\left(\varepsilon_t;\mathbf{x}_t,f_h^{-1}\left(\mathbf{B},\Sigma,\mathbf{Q}\right)\right)\right)}\right) d\left(\varepsilon_1 \ldots \varepsilon_T\right),$$

where $\tilde{\Phi}\left(\mathbf{B},\Sigma,\mathbf{Q},\varepsilon^v\right) = \phi\left(f_h^{-1}\left(\mathbf{B},\Sigma,\mathbf{Q}\right),\varepsilon^v\right)$ and the term υ_{g_h} is called the volume element of the function g_h evaluated at $g_h^{-1}\left(\varepsilon_t;\mathbf{x}_t,f_h^{-1}\left(\mathbf{B},\Sigma,\mathbf{Q}\right)\right)$. Our Equation (11) implies that $\upsilon_{g_h}\left(g_h^{-1}\left(\varepsilon_t;\mathbf{x}_t,f_h^{-1}\left(\mathbf{B},\Sigma,\mathbf{Q}\right)\right)\right) = |\Sigma|^{-\frac{1}{2}}$ for $1 \leq t \leq T$. Hence,

$$\begin{aligned}
&\int \left[\tilde{\Phi}\left(\mathbf{B},\Sigma,\mathbf{Q},\varepsilon^v\right) > \mathbf{0}\right] \left(\prod_{t=1}^{T} \frac{\pi\left(g_h^{-1}\left(\varepsilon_t;\mathbf{x}_t,f_h^{-1}\left(\mathbf{B},\Sigma,\mathbf{Q}\right)\right)|\mathbf{x}_t,\mathbf{B},\Sigma\right)}{\upsilon_{g_h}\left(g_h^{-1}\left(\varepsilon_t;\mathbf{x}_t,f_h^{-1}\left(\mathbf{B},\Sigma,\mathbf{Q}\right)\right)\right)}\right) d\left(\varepsilon_1 \ldots \varepsilon_T\right) \\
&= \int \left[\tilde{\Phi}\left(\mathbf{B},\Sigma,\mathbf{Q},\varepsilon^v\right) > \mathbf{0}\right] \left(\prod_{t=1}^{T} \pi\left(\varepsilon_t\right)\right) d\left(\varepsilon_1 \ldots \varepsilon_T\right) \\
&= \int \left[\tilde{\Phi}\left(\mathbf{B},\Sigma,\mathbf{Q},\varepsilon^v\right) > \mathbf{0}\right] \left(\prod_{s=1}^{v} \pi\left(\varepsilon_{t_s}\right)\right) d\left(\varepsilon_{t_1} \ldots \varepsilon_{t_v}\right). \qquad (16)
\end{aligned}$$

Equation (16) allows us to write the truncated likelihood in Equation (15) as

$$\pi\left(\mathbf{y}^T|\mathbf{B},\Sigma,\mathbf{Q},\Phi\left(\mathbf{B},\Sigma,\mathbf{Q},\mathbf{y}^v,\mathbf{x}^v\right) > \mathbf{0}\right) = \frac{\left[\Phi(\mathbf{B},\Sigma,\mathbf{Q},\mathbf{y}^v,\mathbf{x}^v) > \mathbf{0}\right]\pi\left(\mathbf{y}^T|\mathbf{B},\Sigma\right)}{\omega(\mathbf{B},\Sigma,\mathbf{Q})}, \qquad (17)$$

where $\omega\left(\mathbf{B},\Sigma,\mathbf{Q}\right) = \int \left[\tilde{\Phi}\left(\mathbf{B},\Sigma,\mathbf{Q},\varepsilon^v\right) > \mathbf{0}\right] \left(\prod_{s=1}^{v} \pi\left(\varepsilon_{t_s}\right)\right) d\left(\varepsilon_{t_1} \ldots \varepsilon_{t_v}\right)$. Equation (17) makes clear that the truncated likelihood can be written as a re-weighting of the likelihood function, with weights inversely proportional to the probability of satisfying the restriction.

One would normally choose priors of $(\mathbf{B},\Sigma,\mathbf{Q})$ that are uniform over $O(n)$. When that is the case, $\pi\left(\mathbf{B},\Sigma,\mathbf{Q}\right) = \pi\left(\mathbf{B},\Sigma\right)$, and the posterior of $(\mathbf{B},\Sigma,\mathbf{Q})$ subject to the narrative sign

restrictions is proportional to

$$\pi\left(\mathbf{B},\Sigma,\mathbf{Q}|\mathbf{y}^T,\Phi\left(\mathbf{B},\Sigma,\mathbf{Q},\mathbf{y}^v,\mathbf{x}^v\right)>\mathbf{0}\right)\propto\frac{\left[\Phi\left(\mathbf{B},\Sigma,\mathbf{Q},\mathbf{y}^v,\mathbf{x}^v\right)>\mathbf{0}\right]\pi\left(\mathbf{y}^T|\mathbf{B},\Sigma\right)}{\omega\left(\mathbf{B},\Sigma,\mathbf{Q}\right)}\pi\left(\mathbf{B},\Sigma\right).$$

In other words, the posterior distribution is proportional to the re-weighted likelihood times the prior. On the contrary, as mentioned above, for the case of traditional sign restrictions, it is the prior and not the likelihood that is truncated. Using similar derivations, under priors that are uniform over $O(n)$ the posterior distribution subject to the traditional sign restrictions is $\pi\left(\mathbf{B},\Sigma,\mathbf{Q}|\mathbf{y}^T,\Gamma\left(f_h^{-1}(\mathbf{B},\Sigma,\mathbf{Q})\right)>\mathbf{0}\right)\propto\left[\Gamma\left(f_h^{-1}(\mathbf{B},\Sigma,\mathbf{Q})\right)>\mathbf{0}\right]\pi\left(\mathbf{y}^T|\mathbf{B},\Sigma\right)\pi\left(\mathbf{B},\Sigma\right)$, in which no re-weighting of the likelihood is needed. If one uses both traditional and narrative sign restrictions the posterior distribution $\pi\left(\mathbf{B},\Sigma,\mathbf{Q}|\mathbf{y}^T,\Gamma\left(f_h^{-1}(\mathbf{B},\Sigma,\mathbf{Q})\right)>\mathbf{0},\Phi\left(\mathbf{B},\Sigma,\mathbf{Q},\mathbf{y}^v,\mathbf{x}^v\right)>\mathbf{0}\right)$ is proportional to

$$\left[\Gamma\left(f_h^{-1}(\mathbf{B},\Sigma,\mathbf{Q})\right)>\mathbf{0}\right]\frac{\left[\Phi\left(\mathbf{B},\Sigma,\mathbf{Q},\mathbf{y}^v,\mathbf{x}^v\right)>\mathbf{0}\right]\pi\left(\mathbf{y}^T|\mathbf{B},\Sigma\right)}{\omega\left(\mathbf{B},\Sigma,\mathbf{Q}\right)}\pi\left(\mathbf{B},\Sigma\right).$$

4.2 The algorithm

In practice, one would normally choose priors of $(\mathbf{B},\Sigma,\mathbf{Q})$ that are uniform-normal-inverse-Wishart. In that choice, we are now ready to specify our algorithm to independently draw from the uniform-normal-inverse-Wishart posterior of $(\mathbf{B},\Sigma,\mathbf{Q})$ conditional on the traditional and narrative sign restrictions.

Algorithm 1. *This algorithm makes independent draws from the uniform-normal-inverse-Wishart posterior of* $(\mathbf{B},\Sigma,\mathbf{Q})$ *conditional on the traditional and narrative sign restrictions.*

1. *Independently draw* $(\mathbf{B},\Sigma)$ *from the normal-inverse-Wishart posterior of the reduced-form parameters and* $\mathbf{Q}$ *from the uniform distribution over* $O(n)$.

2. *Check whether* $\left[\Gamma\left(f_h^{-1}(\mathbf{B},\Sigma,\mathbf{Q})\right)>\mathbf{0}\right]$ *and* $\left[\Phi\left(\mathbf{B},\Sigma,\mathbf{Q},\mathbf{y}^v,\mathbf{x}^v\right)>\mathbf{0}\right]$ *are satisfied.*

3. *If not, discard the draw. Otherwise let the importance weight of* $(\mathbf{B}, \Sigma, \mathbf{Q})$ *be as follows*

   3.1. *Simulate M independent draws of* ε^v *from the standard normal distribution.*

   3.2. *Approximate* $\omega(\mathbf{B}, \Sigma, \mathbf{Q})$ *by the proportion of the M draws that satisfy* $\tilde{\Phi}(\mathbf{B}, \Sigma, \mathbf{Q}, \varepsilon^v) > \mathbf{0}$ *and set the importance weight to* $\frac{1}{\omega(\mathbf{B},\Sigma,\mathbf{Q})}$.

4. *Return to Step 1 until the required number of draws has been obtained.*

5. *Draw with replacement from the set of* $(\mathbf{B}, \Sigma, \mathbf{Q})$ *using the importance weights.*

As explained in detail in Arias et al. (2016b), this choice of priors of $(\mathbf{B}, \Sigma, \mathbf{Q})$ is good because it is extremely easy and efficient to make independent draws from the normal-inverse-Wishart distribution and because Rubio-Ramirez et al. (2010) describe how to use the QR decomposition to independently draw the uniform distribution over $O(n)$. Algorithm 1 makes clear that it does not suffice to simply discard the draws that violate the narrative sign restrictions. This would imply giving higher posterior probability to draws of $(\mathbf{B}, \Sigma, \mathbf{Q})$ that are more likely to satisfy the narrative sign restrictions. Hence, this would amount to drawing from a posterior distribution of $(\mathbf{B}, \Sigma, \mathbf{Q})$ that it is not uniform-normal-inverse-Wishart. Instead, we need to compute the importance weights and re-sample the draws accordingly.[4] Also, because of the reasons explained in Arias et al. (2016b), Algorithm 1 is making independent draws from the posterior normal-generalized-normal distribution of $\mathbf{\Theta}$.[5] Further details on the computational properties are provided in Appendix D.

[4]The number of draws M in step 3 needs to be high enough to accurately approximate the importance weights. The larger v is, the more draws will be required. We find that one thousand draws are usually enough to obtain an accurate approximation when narrative restrictions are used in one or two events. For exercises involving more than five or six restrictions, as many as one million might be needed.

[5]See Arias et al. (2016b) for a definition of normal-generalized-normal distribution.

5 Demand and Supply Shocks in the Oil Market

In this section we use narrative information to revisit efforts by Kilian (2009b) and Kilian and Murphy (2012) to assess the relative importance of supply and demand shocks in the oil market. The case of the oil market is particularly well suited for our procedure because a vast literature has documented a number of widely accepted historical events associated with wars or civil conflicts in major oil-producing countries that led to significant physical disruptions in the oil market. We will show that, while the identification scheme proposed by Kilian and Murphy (2012) – based on traditional sign restrictions – does a very good job at separating the effects of supply and demand shocks, adding narrative sign restrictions improves the ability to distinguish between aggregate demand and oil-specific demand shocks, in line with the conclusions of Kilian and Murphy (2014).

5.1 Data and baseline specification

Our starting point is the reduced-form VAR for the global oil market introduced in Kilian (2009b), which has become standard in the literature. The model includes three variables: the growth rate of global oil production, an index of real economic activity, and the log of the real price of oil. To maximize comparability, we choose the exact specification, reduced-form prior and data definitions used in the aforementioned papers.[6]

Kilian and Murphy (2012) and Baumeister and Peersman (2013) use traditional sign restrictions on the contemporaneous IRFs to identify three shocks: an oil supply shock, an aggregate demand shock, and an oil-specific demand shock. In particular, they postulate that impact of these shocks have the signs given in Table 1. Morevoer, Kilian and Murphy (2012)

[6]The VAR is estimated on monthly data using 24 lags and a constant. We extend their data set backward to January 1971 and forward to December 2015. Updated data for the index of real economic activity in Kilian (2009b) were obtained from Lutz Kilian's website, downloaded on March 21, 2016. We refer to the aforementioned papers for details on the sources and the model specification.

Table 1: Sign Restrictions on Impact Responses

Variable \ *Shock*	Oil Supply	Aggregate Demand	Oil-specific Demand
Oil Production	−	+	+
Economic Activity	−	+	−
Real Oil Price	+	+	+

make a compelling argument that many structural parameters that satisfy the sign restrictions in Table 1 imply implausibly large values for the price elasticity of oil supply. This elasticity can be computed from the ratio of the impact responses of production growth and the real price of oil to aggregate demand and oil-specific demand shocks.[7] They propose a plausible upper bound to both of these coefficients of 0.0258, and discard structural parameters that do not satisfy this restriction. We will refer to the traditional sign restrictions formed by Table 1 and the elasticity bounds as the baseline specification.

5.2 The narrative information

We now discuss the narrative information we will use to elicit the narrative sign restrictions. Our main sources are Kilian (2008) and Hamilton (2009), who examined in detail the major exogenous events in the post-1973 period. Figure 1 plots the monthly time series of global oil production growth and the real price of oil, with the following historical events marked as vertical lines: the Yom Kippur War and subsequent Arab oil embargo (October 1973), the start of the Iranian Revolution (December 1978-January 1979), the outbreak of the Iran-Iraq War (September-October 1980), the start of the Persian Gulf War (August 1990), the Venezuela oil strike of December 2002, the start of the Iraq War (March 2003) and the Libyan Civil War (February 2011).[8] It is obvious that these historical events had a major impact both on the

[7]Elasticities are $(\mathbf{L}_0(\mathbf{\Theta}))_{1,2}/(\mathbf{L}_0(\mathbf{\Theta}))_{3,2}$ and $(\mathbf{L}_0(\mathbf{\Theta}))_{1,3}/(\mathbf{L}_0(\mathbf{\Theta}))_{3,3}$, with $(\mathbf{L}_h(\mathbf{\Theta}))_{i,j}$ (i,j) entry of $\mathbf{L}_h(\mathbf{\Theta})$.

[8]The latter event occurred after the publication of the aforementioned papers but there is a good case for including it in the list of exogenous events, see Kilian and Lee (2014). The Libyan Civil War erupted in February

Figure 1: Chronology of Oil Supply Shocks

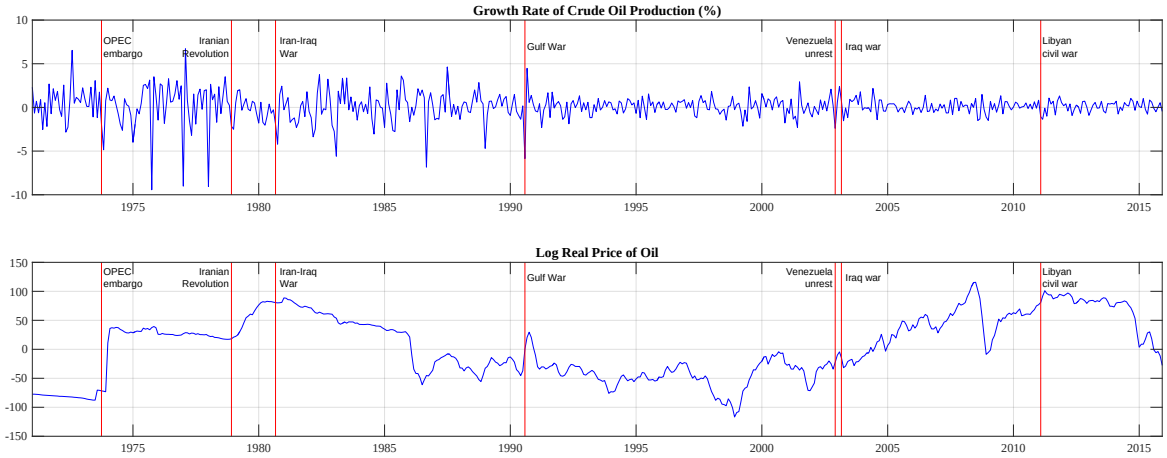

Note: The vertical bars indicate major exogenous oil supply disruptions, associated with the Yom Kippur War and subsequent Arab oil embargo (October 1973), Iranian Revolution (December 1978-January 1979), the Iran-Iraq War (September-October 1980), the Persian Gulf War (August 1990), the Venezuela oil strike of December 2002, the start of the Iraq War (March 2003) and the Libyan Civil War (February 2011).

production growth and the real price of oil. To the extent that these historical events were exogenous with respect to macroeconomic determinants and lowered global oil production, they are a natural candidate for exogenous oil supply shocks.

In any case Barsky and Kilian (2002) and Kilian (2008) have argued against including the 1973 episode in the list of exogenous events, noting that the Arab oil embargo may have been an endogenous response to global demand and US inflationary pressures, and that there is no evidence of OPEC oil production shortfalls having been caused by military action during the Yom Kippur War. Since there is no agreement on this particular event, we exclude the 1973 episode.[9] Thus, we impose the following narrative sign restriction:

2011 in the context of wider protests in favor of civil liberties and human rights in other Arab countries known as the "Arab spring." Before the outbreak of the Civil War, Libya represented over 2% of global crude oil production. From February to April 2011, Libyan production came essentially to a halt.

[9]Moreover, as Kilian (2008) argues, there is a structural change in the oil market around 1973. Prior to 1973 the US price of oil was mostly regulated by government agencies, resulting in extended periods of a constant real price of oil, interrupted only by large discrete jumps. In any case, we have checked the results that will follow, and they are unaffected by adding restrictions based on this event.

Narrative Sign Restriction 1. *The oil supply shock must take negative values in December 1978-January 1979, September-October 1980, August 1990, December 2002, March 2003 and February 2011.*

It is also agreed that the oil supply shocks listed above "resulted in dramatic and immediate disruption of the flow of oil from key global producers" (Hamilton, 2009, p. 220). Therefore, we will use the following narrative sign restriction:

Narrative Sign Restriction 2. *For the periods specified by Restriction 1, oil supply shocks are the most important contributor to the observed unexpected movements in oil production growth. In other words, the absolute value of the contribution of oil supply shocks is larger than the absolute value of the contribution of any other structural shock.*

While Narrative Sign Restriction 2 reflects the agreement that the bulk of the unexpected fall in oil production growth was due to negative oil supply shocks, there is much less agreement in the literature about the ultimate cause of the unexpected increase in the real price of oil. For instance, while Hamilton (2009), p. 224, argues that "oil price shocks of past decades were primarily caused by significant disruptions in crude oil production brought about by largely exogenous geopolitical events," Lutz Kilian, in the comment to the same paper, expresses the view that "a growing body of evidence argues against the notion that the earlier oil price shocks were driven primarily by unexpected disruptions of the global supply of crude oil" (Kilian, 2009a, p. 268.), emphasizing instead the role of the demand for oil. It is possible, however, to find an agreement that "for the oil dates of 1980 and 1990/91 there is no evidence of aggregate demand pressures in industrial commodity markets" (Kilian, 2008, p. 234.). Thus, although there is no agreement on whether oil supply or oil-specific demand shocks caused the unexpected changes in the real price of oil, it seems that both Kilian (2008) and Hamilton (2009) agree that aggregate demand shocks were not responsible for the increases observed in 1980 and 1990. Hence, we will also use the following narrative sign restriction:

Narrative Sign Restriction 3. *For the periods corresponding to September-October 1980 (outbreak of the Iran-Iraq War) and August 1990 (outbreak of the Persian Gulf War), aggregate demand shocks are the least important contributor to the observed unexpected movements in the real price of oil. In other words, the absolute value of the contribution of aggregate demand shocks is smaller than the absolute contribution of any other structural shock.*

In terms of the definitions in Section 3, Narrative Sign Restriction 1 is a restriction on the signs of the structural shocks, whereas Narrative Sign Restrictions 2 and 3 are Type A restrictions on the historical decompositions.

5.3 Results

Figure 2 displays IRFs of the three variables to the three structural shocks, with and without the narrative information. The light shaded area represents the 68% (point-wise) highest posterior density (HPD) credible sets for the IRFs and the dotted lines are the point-wise median IRFs using the baseline identification.[10] The darker shaded areas and solid lines display the equivalent quantities when Narrative Sign Restrictions 1-3 are also used.[11] The narrative sign restrictions dramatically narrow down the uncertainty around many of the IRFs relative to the baseline identification and modifies the shape of some of the IRFs in economically meaningful ways. Oil-specific demand shocks are shown to have a larger contemporaneous effect on the real price of oil that dissipates after around 18 months, whereas aggregate demand shocks have a small initial effect that gradually builds up over time. Some of the IRFs of the economic activity index are also altered substantially. In particular, oil-specific demand shocks have an initial impact on real economic activity that is

[10]It is commonplace to report point-wise median and associated percentiles for the IRFs in the context of set-identified SVAR models. We follow this convention for expository purposes, although Inoue and Kilian (2013), among others, have shown the problems associated with it.

[11]Narrative Sign Restrictions 1-3 affect in total 19 structural shocks. Fifty thousand draws that satisfy the baseline restrictions are generated. Out of these, 920 additionally satisfy the narrative sign restrictions. We approximate their weights in the importance step by using one million draws.

Figure 2: IRFs with and without Narrative Sign Restrictions

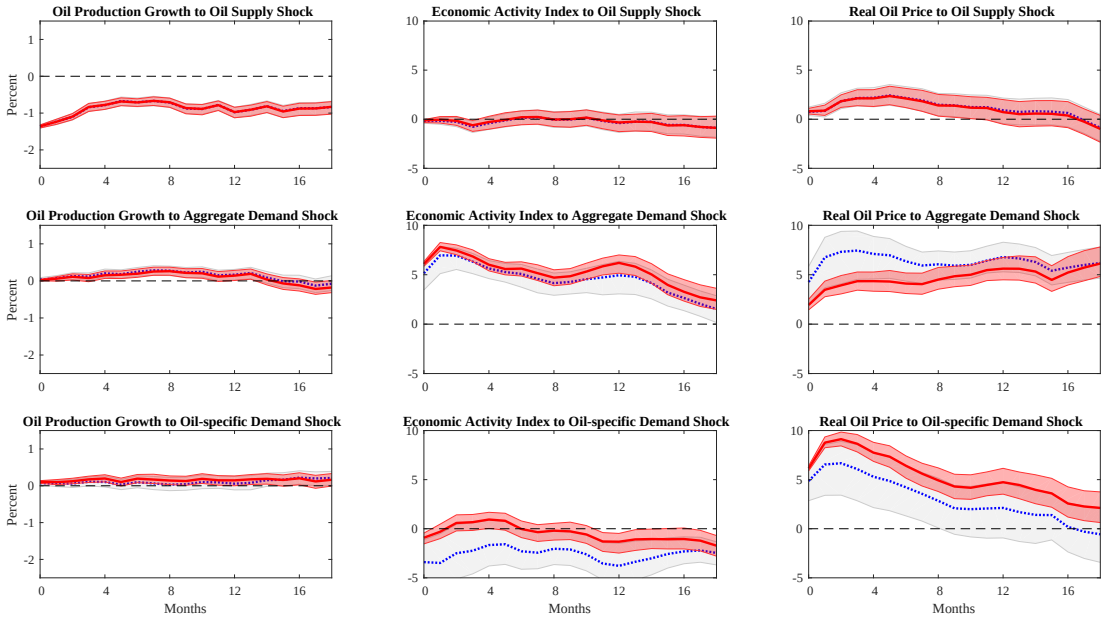

Note: The light shaded area represents the 68% (point-wise) HPD credible sets for the IRFs and the dotted lines are the median IRFs using the baseline identification restrictions. The darker shaded areas and solid lines display the equivalent quantities when Narrative Sign Restrictions 1-3 are also satisfied. Note that the IRF to oil production has been accumulated to the level.

much smaller in absolute value than in the baseline specification. Although it is negative at impact, it builds over time and becomes significant after about 18 months. The response of real economic activity to aggregate demand shocks is stronger and more persistent. The IRFs with the narrative sign restrictions are strikingly similar to the results reported by Kilian (2009b) using the Cholesky decomposition, with the major difference that, in our identification scheme, oil-specific demand shocks are contractionary for economic activity, whereas in the recursive specification these shocks, somewhat counter-intuitively, caused a temporary boom in economic activity.[12]

The economic implications of Narrative Sign Restrictions 1-3 become clear when examining the forecast error variance decompositions (FEVD), which show what fraction

[12]The results using the Cholesky decomposition can be seen in Figure 3, p. 1061, in Kilian (2009b). For the results based on that model specification to make economic sense, the oil-specific demand shock must reflect expectations of rising global aggregate demand for oil.

Figure 3: FEVD with and without Narrative Sign Restrictions

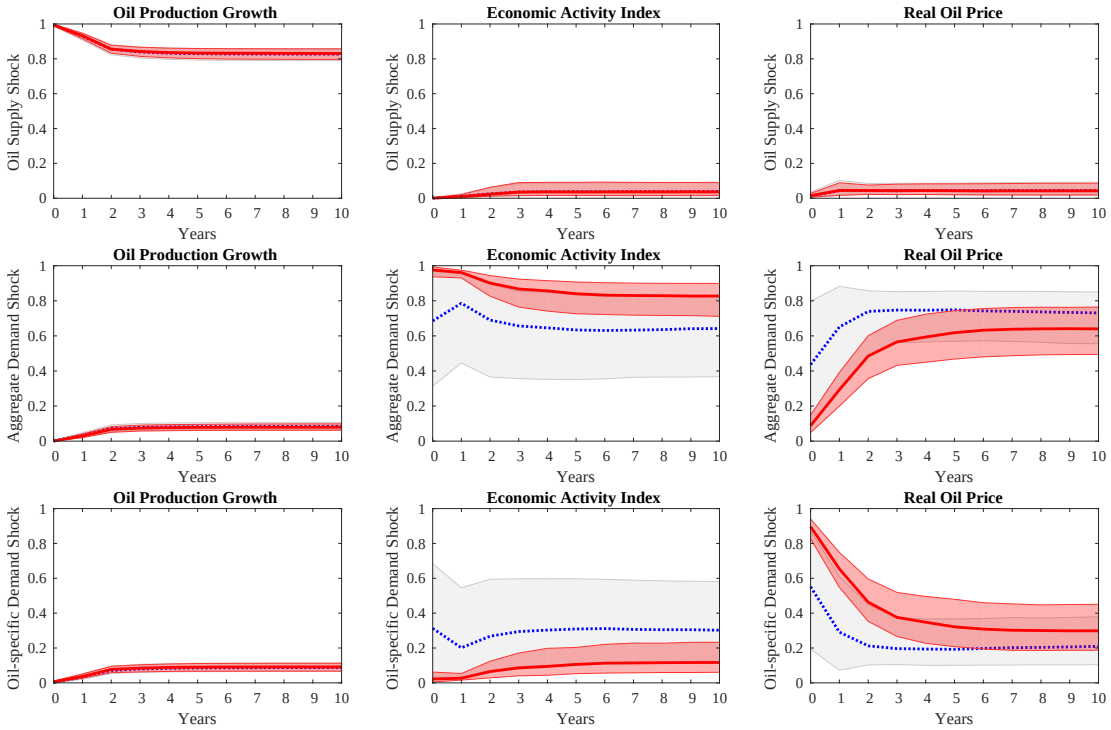

Note: Each panel presents the estimated contribution of each structural shock to the mean squared forecast error at horizons of 1-10 years for the three variables, expressed as a percentage of the total MSE. The light shaded area represents the 68% (point-wise) HPD credible sets for the FEVD, and the dotted lines are the median FEVDs using the baseline identification restrictions. The darker shaded areas and solid lines display the equivalent quantities when Narrative Sign Restrictions 1-3 are also satisfied.

of the unexpected fluctuations in the variables at different horizons can be attributed to each structural shock. Figure 3 shows that when the narrative information and the baseline identification are used, oil-specific demand shocks are responsible for the bulk of the high frequency unexpected variation in the real price of oil, whereas aggregate demand shocks become the most important source of unexpected fluctuations only after three years. With regard to the economic activity index, aggregate demand shocks are now responsible for most of the unexpected fluctuations, although oil supply and oil-specific demand shocks are jointly responsible for over 10% of the unexpected variance in economic activity after ten years. These conclusions clearly contrast with the FEVD obtained using only the baseline

specification, in which oil-specific demand shocks account for about 40% of the unexpected variation in the economic activity index at all horizons and aggregate demand shocks are responsible for the largest share of unexpected fluctuations in the real price of oil even at high frequency. Another important message from Figure 3 is the reduction in uncertainty around the median FEVD. If we compare the light and the darker shaded areas we see that adding the narrative sign restrictions (darker shaded areas) makes the 68% HPD credible sets significantly smaller. Thus, after observing Figures 2 and 3, we can conclude that while the baseline specification, and in particular the restriction on the price elasticity of supply, is very successful at sharpening the effects of oil supply shocks, the narrative information dramatically improves the separate identification of the effects of aggregate demand and oil-specific demand shocks.

To see how the narrative information helps sharpen the identification of aggregate demand and oil-specific demand shocks, it is also informative to examine how Restrictions 1-3 modify the historical decomposition of the real price of oil for particular historical episodes. Panel (a) of Figure 4 looks at the Persian Gulf War, which was one of the events included in Narrative Sign Restrictions 1-3. The baseline identification (light shaded area) is consistent with many structural parameters that imply that aggregate demand shocks were important contributors to the unexpected increase in log real oil prices observed between July and November 1990. Including Narrative Sign Restrictions 1-3 (darker shaded area) reinforces the view of Kilian and Murphy (2014) that speculation in the physical market, i.e., an oil-specific demand shock, was the cause of the bulk of the unexpected 60% increase in the real price of oil at the outbreak of the war. Panels (b) and (c) look at two events for which no restrictions are imposed. For the run-up in the real price of oil between 2004-2008, displayed in Panel (b), the narrative information agrees with the baseline identification in that aggregate demand shocks were the main cause. This is in line with the results of the previous literature. For the 60% unexpected decline in the real price of oil observed between July 2014 and December

Figure 4: Historical Decomposition of Oil Price Movements around Selected Episodes

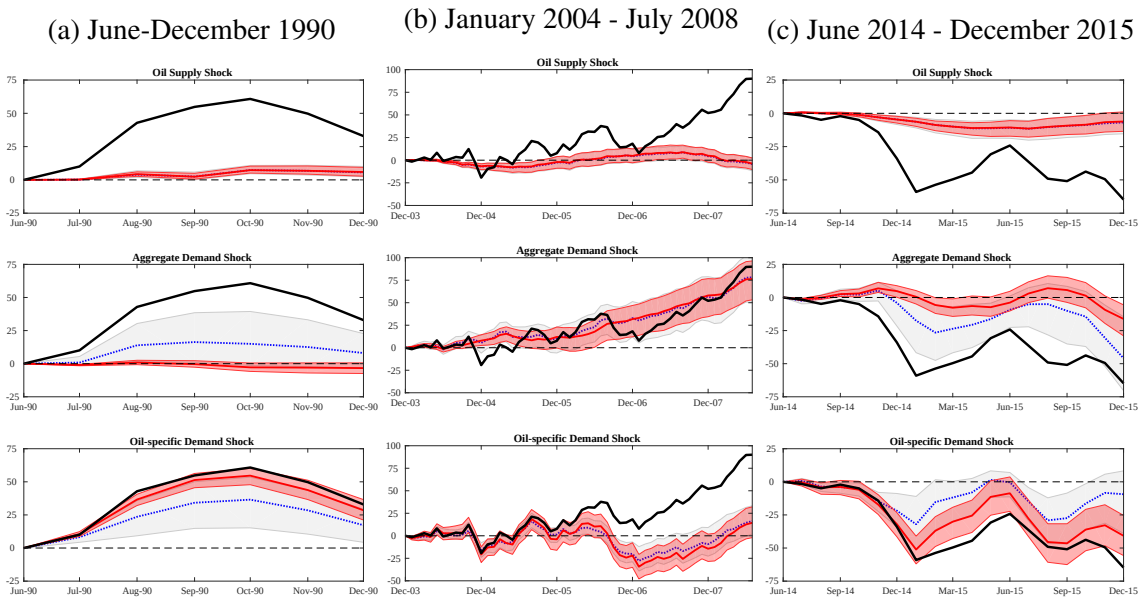

Note: For selected historical episodes, the panels display the observed unexpected change in the real price of oil (in log points) attributed to each of the structural shocks. The observed unexpected change is represented by the solid thick line. The dotted lines are the median for the baseline identification restrictions, while the light shaded area represents the 68% (point-wise) HPD credible sets. The solid thin lines and the darker shaded areas display the equivalent quantities when Narrative Sign Restrictions 1-3 are also satisfied.

2015, Panel (c) shows how the baseline identification concludes that it was not due to oil supply shocks, but leads to substantial uncertainty about whether aggregate demand shocks or oil-specific demand shocks were behind the collapse. With the narrative information the results point toward oil-specific demand shocks as the source of the collapse.

5.4 Assessing the importance of each historical event

Because we focus on a small number of historical events, it is straightforward to assess the importance of each of them. Table 2 computes what percentage of draws of the structural parameters that satisfy the baseline specification violates each of the narrative sign restrictions, both individually and jointly. It is important to note that a high probability of violating

Table 2: Probability of Violating the Narrative Sign Restrictions

	Restr. 1	Restr. 2	Restr. 3	Any Restr.
Iranian Revolution	20%	2.9%	–	21%
Iran-Iraq War	0%	0%	46%	46%
Gulf War	0%	0%	93%	93%
Venezuela Unrest	0%	0%	–	0%
Iraq War	43%	21%	–	53%
Libyan Civil War	4.6%	1%	–	5%
Any Episodes	42%	24%	93%	98%

a restriction should not be interpreted as evidence against its validity. On the contrary, it tells us that the baseline specification admits many structural parameters that, according to the narrative sign restrictions, should be rejected. Therefore, the higher the probability of violating a narrative sign restriction, the more informative the restriction will be for achieving identification.[13] The results indicate that Narrative Sign Restrictions 1 and 2 are less relevant than Narrative Sign Restriction 3. However, it is noteworthy that the baseline identification still includes many structural parameters for which a positive supply shock occurred during either the 1979 Iranian Revolution or the 2003 Iraq War, contradicting Narrative Sign Restriction 1. In total, 42% of the structural parameters that satisfy the baseline specification violate Narrative Sign Restriction 1. It is also the case that over 20% of the structural parameters that satisfy the baseline specification do not satisfy Narrative Sign Restriction 2 for the 1979 Iranian Revolution or the 2003 Iraq War. But it is clear that Narrative Sign Restriction 3 is key to obtaining the results of Figures 2 and 3, given that in total 93% of the structural parameters that satisfy the baseline specification do not respect Narrative Sign Restriction 3.

In fact, it turns out that to obtain the results of Figures 2 and 3 it is sufficient to impose Restriction 3 for the August 1990 event. Using an Alternative Narrative Sign Restriction 3

[13]For a similar point, see Kilian and Lütkepohl, (2017, chapter 13)

that includes only this event leads to results that are indistinguishable from the ones presented above.[14] In other words, one only needs to agree that expansionary aggregate demand shocks were the least important contributor to the unexpected spike in the real price of oil observed that month, a view that has been described as agreeable to a wide group of experts (Kilian and Murphy, 2014, p. 468), to obtain our results.

Given that the challenge is to come up with additional uncontentious sign restrictions that help shrink the set of admissible structural parameters, the resemblance of the results using either Narrative Sign Restrictions 1-3 or Alternative Narrative Sign Restriction 3 is a great success. By using a single narrative sign restriction to constraint the set of structural parameters to those whose implied behavior in August 1990 agrees with the generally accepted description of that event, we can greatly sharpen the separate identification of aggregate demand and oil-specific demand shocks for the entire sample, including many other periods for which narrative information is not available.

6 Monetary Policy Shocks and the Volcker Reform

An extensive literature has studied the effect of monetary policy shocks on output using SVARs, identified with zero restrictions, as in Christiano et al. (1999), Bernanke and Mihov (1998), sign restrictions, as in Uhlig (2005), or both, as in Arias et al. (2016a). SVARs identified using zero restrictions have consistently found that an exogenous increase in the fed funds rate induces a reduction in real activity. This intuitive result has become the "consensus." This consensus view, however, has been challenged by Uhlig (2005), who criticizes imposing a questionable zero restriction on the IRF of output to a monetary policy shock on impact. To solve the problem he proposes to identify a shock to monetary policy by imposing sign restrictions only on the IRFs of prices and nonborrowed reserves to this shock,

[14]For this reason the figures are omitted here, but this result is presented in the Appendix A.

while imposing no restrictions on the IRF of output. The lack of restrictions on the IRF of output to a monetary policy shock makes this is an attractive approach. Importantly, under his identification, an exogenous unexpected increase in the fed funds rate does not necessarily induce a reduction in real activity.

An alternative approach to identify the effects of monetary policy shocks uses historical sources to isolate events that constitute exogenous monetary policy shocks. Following the pioneering work of Friedman and Schwartz (1963), Romer and Romer (1989) combed through the minutes of the FOMC to create a dummy time series of events that they argued represented exogenous tightenings of monetary policy. Focusing exclusively on contractionary shocks, they singled out a handful of episodes in the postwar period "in which the Federal Reserve attempted to exert a contractionary influence on the economy in order to reduce inflation" (Romer and Romer (1989) , p. 134). The Romers' monetary policy time series narrative has become very influential, but has been criticized by Leeper (1997), who pointed out that their dates are predictable from past macroeconomic data. As a consequence, in recent years alternative methods have been developed to construct time series of monetary policy shocks that are by design exogenous to the information set available at the time of the policy decision. The first prominent example is Romer and Romer (2004), who regressed changes of the intended federal funds rate between FOMC meetings on changes in the Fed's Greenbook forecasts of output and inflation. By construction, the residuals from this regression are orthogonal to all the information contained in the Greenbook forecasts, and can plausibly be taken to be a measure of exogenous monetary policy shocks. A second approach looks at high-frequency financial data. Kuttner (2001), Gürkaynak et al. (2005), and Gertler and Karadi (2015) look at movements in federal funds futures contracts during a short window around the time of policy announcements to isolate the monetary policy shocks.

However, the existing narrative time series are sometimes inconclusive and other times contradictory. This is not just due to differences in methods and sources, but, as Ramey

(2016) recently pointed out, to the fact that the Federal Reserve has historically reacted in a systematic way to output and inflation developments (see also Leeper et al., 1996). This systematic response is a key difference with the oil supply shocks analyzed in Section 5, so the occurrence and importance of truly exogenous monetary policy shocks remain controversial. Thus, monetary policy shocks are much more difficult to isolate than oil supply shocks.

For this reason, in this section we will use narrative sign restrictions for a single event: October of 1979. The monetary policy decisions of October 6, 1979, enacted shortly after Paul Volcker became chairman of the Fed, are described by Romer and Romer (1989) as "a major anti-inflationary shock to monetary policy" and represent, in our view, the clearest case in the postwar period of an exogenous monetary policy shock. Lindsey et al. (2013) provide a detailed account of the events leading to the decision to abandon targeting the federal funds rate in favor of targeting non borrowed reserves as the operating procedure for controlling the money supply. While macroeconomic conditions, in particular, the deterioration of the inflation outlook and the increase in the real price of oil that followed the Iranian Revolution of 1978-79, played a large role in causing the shift, the forcefulness and the surprise character of the action and the dramatic break with established practice in the conduct of policy strongly suggest the occurrence of a monetary policy shock.

An argument could be made that the Volcker event is so large that it may be best modeled as a change in the monetary rule rather than a monetary policy shock. But Sims and Zha (2006) contend that the evidence for changes in the parameters of the monetary policy rule during the Volcker period is weak and that this period is best described as a period of high variance in the monetary policy shocks to an otherwise unchanged monetary policy rule. Primiceri (2005) reaches similar conclusions. Indeed, as Lindsey et al. (2013) describe, contemporaneous observers accused the FOMC of adopting the new operating procedures only as a smokescreen to obscure its intention to markedly increase short-term interest rates. We side with this view and regard the episode as a large shock within a stable monetary rule.

6.1 Data and baseline specification

Our starting point is the reduced-form VAR used, among others, by Christiano et al. (1999), Bernanke and Mihov (1998) and Uhlig (2005). The model includes six variables: real output, the GDP deflator, a commodity price index, total reserves, nonborrowed reserves, and the federal funds rate. As in the previous section, we chose the exact specification, reduced-form prior and data definitions used in the aforementioned papers. Our sample period is January 1965 to November 2007.[15] Our baseline identification is identical to Uhlig's (2005). Specifically, he postulates that a contractionary monetary policy shock increases the federal funds rate and reduces the GDP deflator, the commodity price index and non-borrowed reserves for periods 0 to 5 months.

6.2 The narrative information

We start by examining the implications of the baseline specification for the period around October 1979. The light histogram in Panel (a) of Figure 5 displays the posterior distribution of the monetary policy shock during that month. While most of the distribution has positive support (i.e., a contractionary monetary policy shock occurred), the baseline identification implies that a negative (i.e., expansionary) monetary policy shock occurred with about an 11% posterior probability. Panel (b) plots the counterfactual path (dotted line with light 68% point-wise HPD credible sets) of the federal funds rate if no structural shock other than the monetary policy shock had occurred between September 1979 and December 1980. As can be seen from Panel (b), the baseline specification implies that the monetary policy shock was rather unimportant in explaining the unexpected increase in the federal funds rate observed in

[15]The VAR uses monthly data using 12 lags and no constant or deterministic trends. We refer to the aforementioned papers for details on the model specification. Following Arias et al. (2016a), we stop the sample in November 2007 because starting in December 2007 there are large movements in reserves associated with the global financial crisis. Furthermore, the federal funds rate has been at the zero lower bound since November 2008. Including the post-crisis sample could obscure the comparison with the results of earlier papers.

Figure 5: Results Around October 1979 with and without Narrative Restrictions

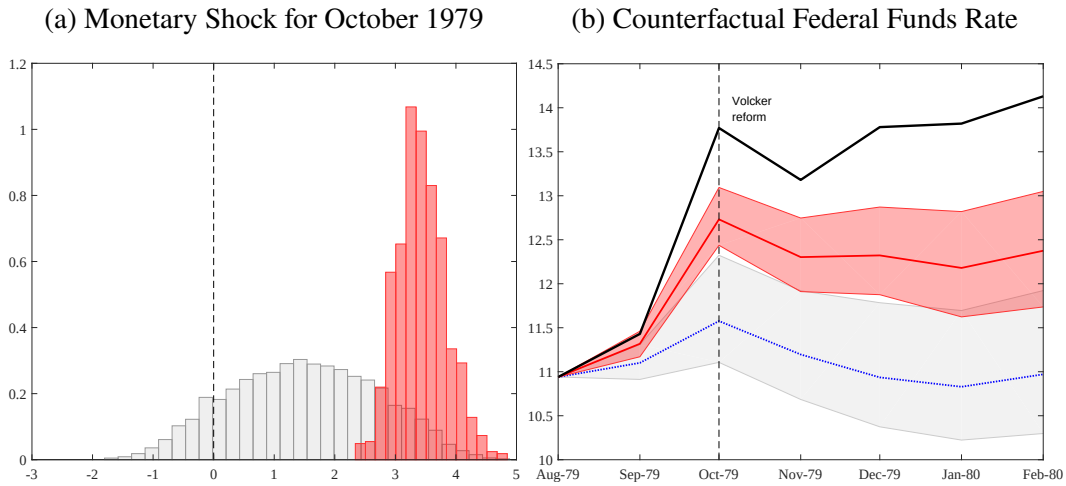

Note: The light histogram in Panel (a) plots the posterior distribution of the monetary policy shock for October 1979 using the baseline identification. The darker histogram plots the same distribution after incorporating Narrative Restrictions 4-5. Panel (b) plots the actual federal funds rate (solid thick) and the median of the counterfactual federal funds rate (dotted) resulting from excluding all non-monetary structural shocks using the baseline identification. The light bands represent 68% (point-wise) HPD credible sets around the median. The solid thin line and darker shaded area plots the same result using narrative restrictions 4-5.

October. So the baseline specification effectively implies that the increase in the federal funds rate between September 1979 and December 1980 was due to some structural shock other than the monetary policy shock. This means that the set of admissible structural parameters implied by the baseline identification retains many structural parameters that go against the widely shared view that in October of 1979 a major contractionary monetary policy shock greatly increased the fed funds rate. In order to eliminate such structural parameters, we therefore impose the following two narrative sign restrictions:

Narrative Sign Restriction 4. *The monetary policy shock for the observation corresponding to October 1979 must be of positive value.*

Narrative Sign Restriction 5. *For the observation corresponding to October 1979, a monetary policy shock is the overwhelming driver of the unexpected movement in the federal*

Figure 6: IRFs with and without Narrative Sign Restrictions

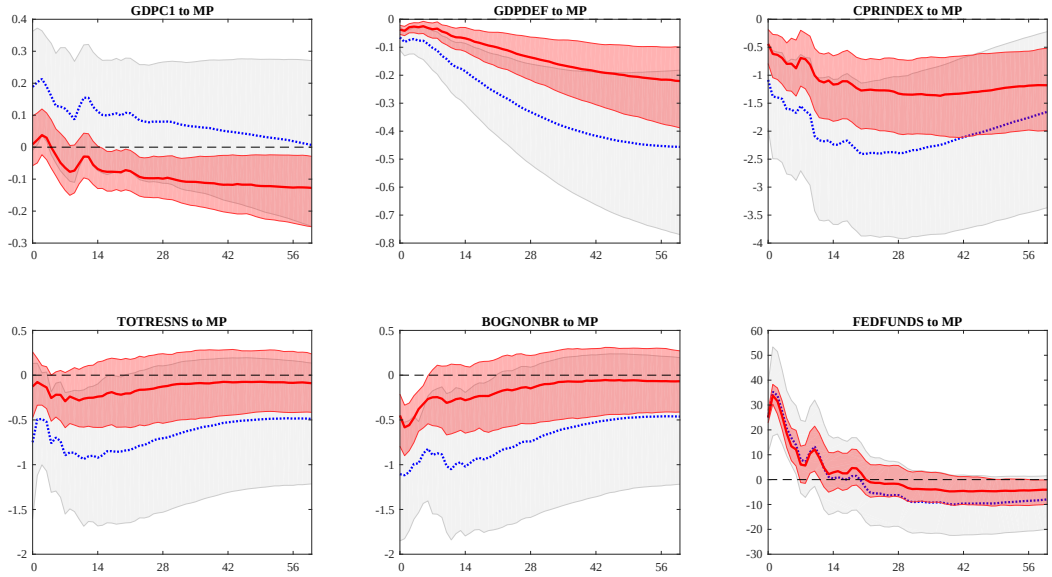

Note: The light shaded area represents the 68% (point-wise) HPD credible sets for the IRFs and the dotted lines are the median IRFs using the baseline identification restrictions. The darker shaded areas and solid lines display the equivalent quantities for the models that additionally satisfy Narrative Sign Restrictions 4 and 5. The monetary policy shock has been normalized to have an impact of 25 basis points on the federal funds rate.

funds rate. In other words, the absolute value of the contribution of monetary policy shocks to the unexpected movement in the federal funds rate is larger than the sum of the absolute value of the contributions of all other structural shocks.

Importantly, we do not place any restrictions on the contribution of the monetary policy shock to the unexpected change in output during that episode, but just on its contribution to the unexpected movement in the federal funds rate. In terms of the definitions of Section 3, Narrative Sign Restriction 4 is a restriction on the sign of the structural shock, whereas Narrative Sign Restriction 5 is a Type B restriction on the historical decomposition.

6.3 Results

Figure 6 compares the IRFs to a monetary policy shock, with and without narrative sign restrictions. The light shaded area represents the 68% (point-wise) HPD credible sets for

the IRFs and the dotted lines are the median IRFs using the baseline identification. These results replicate the IRFs depicted in Figure 6 of Uhlig (2005). The darker shaded areas and solid lines display the equivalent quantities when Narrative Sign Restrictions 4 and 5 are also used.[16] As one can observe, the inclusion of narrative sign restrictions is enough to imply that contractionary monetary policy shocks cause output to drop with very high posterior probability. The results reported highlight that the narrative information embedded in a single event can shrink the set of admissible structural parameters so dramatically that the economic implications change.

How do Narrative Sign Restrictions 4 and 5 change the implications for the period around October 1979? The darker histogram in Panel (a) of Figure 5 displays the posterior distribution of the monetary policy shock during that month when Narrative Sign Restrictions 4 and 5 are also used. The distribution of the structural shock now has positive support with 100% probability. Panel (b) plots the counterfactual path (solid thin line with darker 68% point-wise HPD credible sets) of the federal funds rate, as described above. The monetary policy shock was the overwhelming contributor to the unexpected increase in the federal funds rate. The results indicate that the monetary policy shock was very large (between 2 and 5 standard deviations) and that it was responsible for between 100 and 150 basis points of the roughly 225-basis-point unexpected increase in the federal funds rate observed in October 1979. It is important to emphasize that these magnitudes are not imposed by Narrative Sign Restrictions 4 and 5; only the sign of the shock and the sign of the contribution of the monetary policy shock relative to other structural shocks are.[17] Therefore, if one agrees with the baseline restrictions and also with the fact that the monetary policy shock was both positive and the most important contributor to the October 1979 tightening, one should

[16]Narrative Sign Restrictions 4 and 5 affect in total one structural shock. We obtain 10,116 draws that satisfy the baseline restrictions. Out of these, 931 additionally satisfy Narrative Sign Restrictions 4 and 5. We approximate their weights in the importance step by using one thousand draws.

[17]Results with Type A Narrative Sign Restriction 5 are similar and are available in Appendix B.

conclude that monetary policy shocks reduce output with a high probability.[18]

6.4 Including additional events

The results above have highlighted that using narrative information for the October 1979 event is highly informative. That event is in our view the clearest and most uncontroversial example of a monetary policy shock. However, an important caveat must follow. As Romer and Romer (1989, p. 123) recognize, "the narrative identification of shocks generally occurs retrospectively, and thus [...] there may be an unconscious bias toward, for example, searching harder for negative monetary shocks in periods preceding sharp declines in money and output than in other periods." This was indeed the case for the Volcker event, which was followed by a deep recession. Therefore, it is possible that the proximity of the Volcker announcement to the recession has shaped the historical interpretation of this period as a clear case study of a monetary policy shock.

Fortunately, there is a long literature that uses historical sources to isolate monetary policy shocks and that can provide additional case studies. Owing to differences in sources and methods, the existing narrative series are sometimes inconclusive and other times contradictory. However, by cross-checking the original Romer and Romer (1989) chronology, the updated Greenbook residual series from Romer and Romer (2004), the high-frequency series from Gürkaynak et al. (2005), and the transcripts from the meetings of the FOMC, we have selected seven additional events for which there appears to be reasonable agreement that an important monetary policy shock occurred.[19] Of these, three – April 1974, December 1988 and February 1994 – were contractionary shocks and four – December 1990, October 1998, April 2001, and November 2002 – were expansionary shocks. The February 1994 event is of particular interest because the historical record identifies a major monetary policy

[18]Our results echo those of Inoue and Kilian (2013) and Arias et al. (2016a), which question the robustness of Uhlig's (2005) results, and call for the introduction of additional restrictions.

[19]Details of how we decided to choose those seven events are available in Appendix C.

shock, but output accelerated during 1994.[20] Thus, it is arguably not subject to the criticism that the historical narrative might have been shaped by the presence of a recession.

Once we have selected those additional seven events, it seems reasonable to ask whether adding them as narrative sign restrictions changes the results reported above. For that purpose, we can therefore consider the following narrative sign restrictions:

Narrative Sign Restriction 6. *The monetary policy shock must be positive for the observations corresponding to April 1974, October 1979, December 1988, and February 1994, and negative for December 1990, October 1998, April 2001, and November 2002.*

Narrative Sign Restriction 7. *For the periods specified by Restriction 6, monetary policy shocks are the most important contributor to the observed unexpected movements in the federal funds rate. In other words, the absolute value of the contribution of monetary policy shocks is larger than the absolute value of the contribution of any other structural shock.*

The results using Narrative Sign Restriction 6 and 7 are very similar to those using only October 1979, albeit with narrower HPD credible sets because of additional information.[21] In addition, it is possible to obtain equally similar results by just imposing narrative sign restrictions for the December 1988, February 1994 or April 2001 dates on their own.[22] Hence, we can claim that, while the 1979 event is sufficient to obtain the results reported in Figures 5 and 6, it is not necessary. This leads us to conclude that for the current application, the informativeness of narrative sign restrictions is not dependent on the particularities of the October 1979 episode. Moreover, the fact that the results can be obtained by imposing the 1994 event on its own is particularly reassuring, since it means that the results do not depend

[20]See "The great bond massacre" (Fortune, 1994) for a representative contemporary account, which associated the heavy losses experienced by financial companies, hedge funds, and bond mutual funds on their holdings of long-term bonds with the surprise tightening by the Fed.

[21]Note that Narrative Sign Restriction 7 is a Type A restriction on the historical decomposition, whereas Narrative Sign Restriction 5 is a Type B restriction on the historical decomposition, hence the October 1979 event can be imposed in its less strong version when other events are included.

[22]Owing to space considerations, the figures are omitted here, but this result is presented in Appendix C.

on using events that were followed by a recession.

7 Conclusion

Historical sources have long been regarded as useful for identifying structural shocks. In this paper, we have shown how to use narrative sign restrictions to identify SVARs. We place sign restrictions on structural shocks and the historical decomposition of the data at certain historical periods, ensuring that the structural parameters are consistent with the established narrative account of these episodes. We have illustrated our approach with the case of oil and monetary shocks. We have shown that adding a small number of narrative sign restrictions related to key historical events, and sometimes even a single event, can dramatically sharpen the inference or even alter the conclusions of SVARs only identified with traditional sign restrictions. Relative to existing narrative information methods, our approach has the advantage of requiring that we trust only the sign and the relative importance of the structural shock for a small number of events, which facilitates the practice of basing inference on a few uncontroversial sign restrictions on which the majority of researchers agree and which lead to robust conclusions.

References

Arias, J. E., D. Caldara, and J. F. Rubio-Ramirez (2016a): "The Systematic Component of Monetary Policy in SVARs: An Agnostic Identification Procedure," International Finance Discussion Papers 1131, Board of Governors of the Federal Reserve System (U.S.).

Arias, J. E., J. F. Rubio-Ramirez, and D. F. Waggoner (2016b): "Inference Based on SVARs Identified with Sign and Zero Restrictions: Theory and Applications," Working paper, Federal Reserve Bank of Atlanta.

Barsky, R. B. and L. Kilian (2002): "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," in *NBER Macroeconomics Annual 2001, Volume 16*, National Bureau of Economic Research, Inc, NBER Chapters, 137–198.

Baumeister, C. and G. Peersman (2013): "The role of time-varying price elasticities in accounting for volatility changes in the crude oil market," *Journal of Applied Econometrics*, 28, 1087–1109.

Bernanke, B. S. and I. Mihov (1998): "Measuring Monetary Policy," *The Quarterly Journal of Economics*, 113, 869–902.

Canova, F. and G. D. Nicolo (2002): "Monetary disturbances matter for business fluctuations in the G-7," *Journal of Monetary Economics*, 49, 1131–1159.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999): "Monetary policy shocks: What have we learned and to what end?" in *Handbook of Macroeconomics*, ed. by J. B. Taylor and M. Woodford, Elsevier, vol. 1 of *Handbook of Macroeconomics*, chap. 2, 65–148.

Faust, J. (1998): "The robustness of identified VAR conclusions about money," International Finance Discussion Papers 610, Board of Governors of the Federal Reserve System (U.S.).

Friedman, M. and A. Schwartz (1963): *A Monetary History of the United States, 1867-1960*, National Bureau of Economic Research, Inc.

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

Gürkaynak, R. S., B. Sack, and E. Swanson (2005): "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*, 1.

Hamilton, J. D. (1985): "Historical Causes of Postwar Oil Shocks and Recessions," *The Energy Journal*, 0, 97–116.

——— (2009): "Causes and Consequences of the Oil Shock of 2007-08," *Brookings Papers on Economic Activity*, 40, 215–283.

Inoue, A. and L. Kilian (2013): "Inference on impulse response functions in structural VAR models," *Journal of Econometrics*, 177, 1–13.

Kilian, L. (2008): "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?" *The Review of Economics and Statistics*, 90, 216–240.

——— (2009a): "Comment on 'Causes and Consequences of the Oil Shock of 2007-08', by James D. Hamilton," *Brookings Papers on Economic Activity*, 1 (Spring), 267–278.

——— (2009b): "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 99, 1053–69.

Kilian, L. and T. K. Lee (2014): "Quantifying the speculative component in the real price of oil: The role of global oil inventories," *Journal of International Money and Finance*, 42, 71–87.

Kilian, L. and H. Lütkepohl (2017): *Structural Vector Autoregressive Analysis*, Themes in Modern Econometrics, 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.

——— (2014): "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil," *Journal of Applied Econometrics*, 29, 454–478.

Kuttner, K. N. (2001): "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market," *Journal of Monetary Economics*, 47, 523–544.

Leeper, E. M. (1997): "Narrative and VAR approaches to monetary policy: Common identification problems," *Journal of Monetary Economics*, 40, 641–657.

Leeper, E. M., C. A. Sims, and T. Zha (1996): "What does monetary policy do?" *Brookings Papers on Economic Activity*, 1996, 1–78.

Lindsey, D. E., A. Orphanides, and R. H. Rasche (2013): "The Reform of October 1979: How It Happened and Why," *Federal Reserve Bank of St. Louis Review*, 487–542.

Mertens, K. and M. O. Ravn (2013): "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States," *American Economic Review*, 103, 1212–47.

Montiel-Olea, J. L., J. H. Stock, and M. W. Watson (2015): "Uniform Inference in Structural VARs Identified with External Instruments," *Working Paper*.

Primiceri, G. E. (2005): "Time Varying Structural Vector Autoregressions and Monetary Policy," *Review of Economic Studies*, 72, 821–852.

Ramey, V. A. (2011): "Identifying Government Spending Shocks: It's all in the Timing," *The Quarterly Journal of Economics*, 126, 1–50.

——— (2016): "Macroeconomic Shocks and Their Propagation," NBER Working Papers 21978, National Bureau of Economic Research, Inc.

Romer, C. D. and D. H. Romer (1989): "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz," in *NBER Macroeconomics Annual 1989, Volume 4*, National Bureau of Economic Research, Inc, NBER Chapters, 121–184.

——— (1994): "Monetary policy matters," *Journal of Monetary Economics*, 34, 75–88.

——— (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 94, 1055–1084.

——— (2010): "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks," *American Economic Review*, 100, 763–801.

Rubio-Ramirez, J. F., D. F. Waggoner, and T. Zha (2010): "Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference," *Review of Economic Studies*, 77, 665–696.

Sims, C. A. and T. Zha (2006): "Were There Regime Switches in U.S. Monetary Policy?" *American Economic Review*, 96, 54–81.

Stock, J. and M. Watson (2008): "NBER Summer Institute Minicourse 2008: What's New in Econometrics - Time Series, Lecture 7: Structural VARs," Cambridge, Mass. National Bureau of Economic Research.

——— (2012): "Disentangling the Channels of the 2007-09 Recession," *Brookings Papers on Economic Activity*, Spring, 81–156.

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

Not-for-publication Appendix to "Narrative Sign Restrictions for SVARs"

by Juan Antolin-Diaz and Juan F. Rubio-Ramirez

A Robustness of Results for Oil Market

Consider the following alternative formulation of Narrative Sign Restriction 3.

Alternative Narrative Sign Restriction 3. *For the period corresponding to August 1990 (outbreak of the Persian Gulf War), aggregate demand shocks are the least important contributor to the observed unexpected movements in the real price of oil. In other words, the absolute value of the contribution of aggregate demand shocks is smaller than the absolute value of the contribution of any other structural shock.*

Figure A.1 plots the same IRFs reported in Figure 2, but the darker shaded areas and solid lines now add the Alternative Narrative Sign Restriction 3 to the baseline identification instead of adding the Narrative Sign Restrictions 1-3.[1] As the reader can see, Figures 2 and A.1 are almost identical.[2] Hence using either set of narrative sign restrictions has comparable effects on the IRFs and on other results such as the FEVD and historical decompositions presented above.[3]

Given that the restriction relating to August 1990 appears to be key to our results, it warrants some additional discussion. In particular, we will analyze the robustness of the

[1] Alternatively, one may also reformulate Narrative Sign Restrictions 1 and 2 so as to include only the August 1990 event, but as can be seen from the third row of Table 2, Narrative Sign Restrictions 1 and 2 are always satisfied by the baseline specification for this particular event. Therefore it is enough to use just Alternative Narrative Sign Restriction 3.

[2] Alternative Narrative Sign Restriction 3 affects in total one time period. Ten thousand draws that satisfy the baseline restrictions are generated. Out of these, 749 satisfy the narrative sign restrictions. We approximate their weight in the importance step by using one thousand draws.

[3] The equivalents to Figures 3 and 4 using Alternative Narrative Sign Restriction 3 are essentially identical to the originals, which use Restrictions 1-3. We do not display them owing to space considerations, but they are available upon request.

Figure A.1: IRFs with and without Narrative Sign Restrictions

(Alternative Narrative Sign Restriction 3)

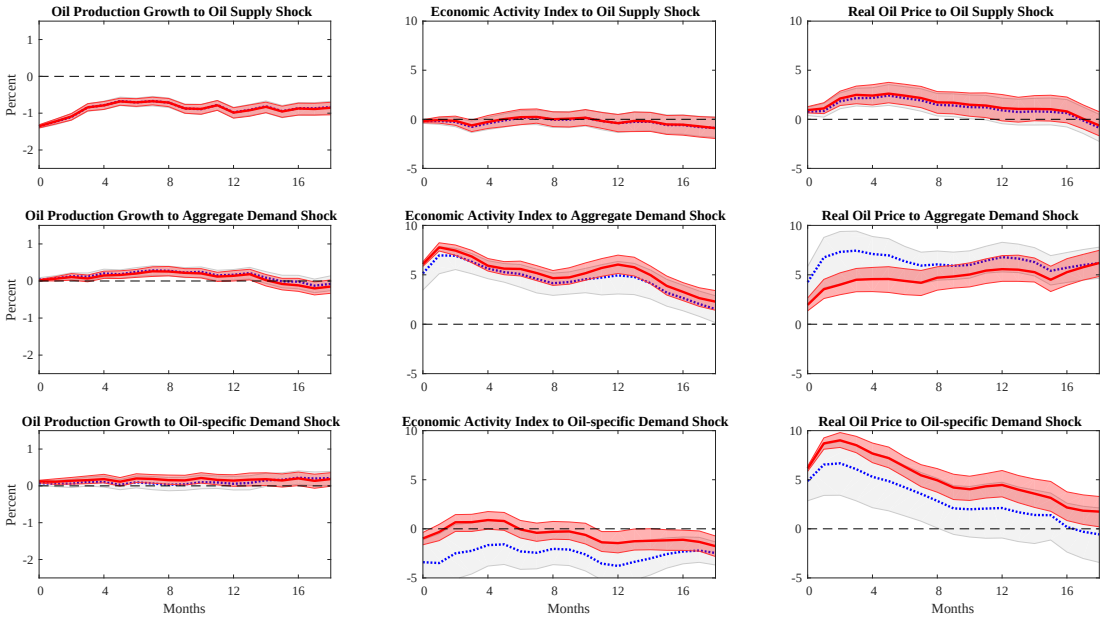

Note: The light shaded area represents the 68% (point-wise) HPD credible sets for the IRFs and the dotted lines are the median IRFs using the baseline identification restrictions. The darker shaded areas and solid lines display the equivalent quantities when Alternative Narrative Sign Restriction 3 is also satisfied. Note that the IRF to oil production has been accumulated to the level.

results to using the Type B variant of Alternative Narrative Sign Restriction 3, instead of the Type A variant we have been using so far. Recall from Section 3.3 that for this case the Type A restriction specifies that the contribution of the aggregate demand shock to the spike in the real price of oil is "less important than any other," whereas the Type B restriction would specify that the contribution is "less important than the sum of all others." Clearly, in this case Type A is a stronger version than Type B, since being less important than any other contribution automatically implies being less important than the (absolute) sum of all others (see the discussion in Section 3.4.2). Figure A.2 plots the same IRFs reported in Figure A.1 when adding the Alternative Narrative Sign Restriction 3 to the baseline identification, but in its milder Type B variant.[4] As the reader can see, the main conclusions are maintained.

[4]Alternative Narrative Sign Restriction 3 (Type B) affects in total two structural shocks. Ten thousand

Figure A.2: IRFs with and without Sign Narrative Restrictions

(Alternative Narrative Sign Restriction 3 – Type B)

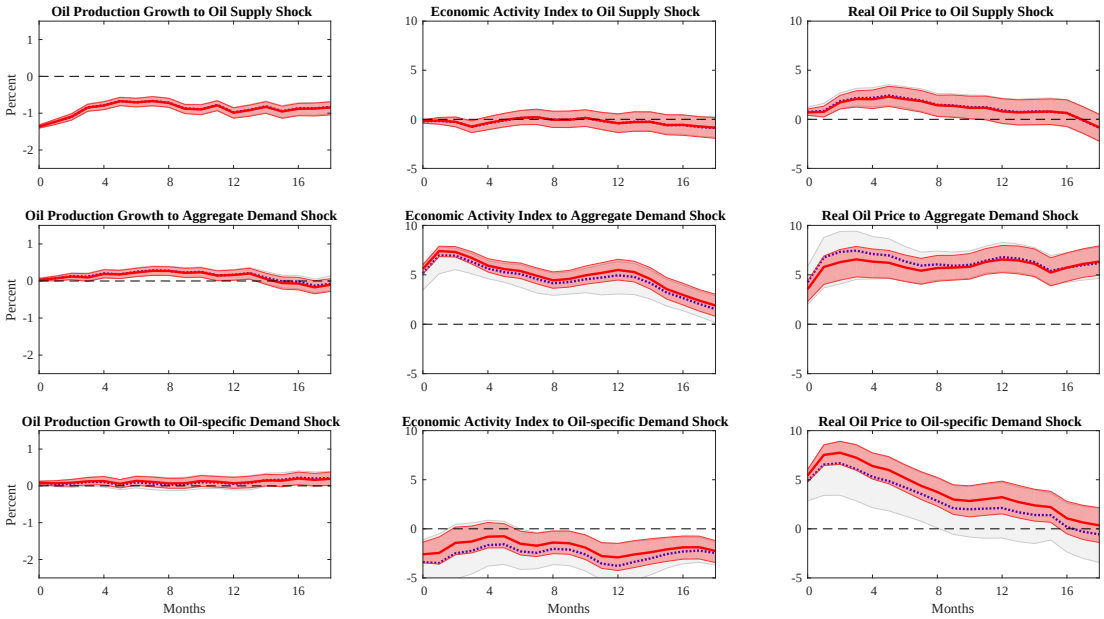

Note: The light shaded area represents the 68% (point-wise) HPD credible sets for the IRFs and the dotted lines are the median IRFs using the baseline identification restrictions. The darker shaded areas and solid lines display the equivalent quantities when the Alternative Narrative Sign Restriction 3 (Type B) is also satisfied. Note that the IRF to oil production has been accumulated to the level.

In any case, since it seems accepted that aggregate demand shocks are the least important contributor to the observed unexpected movements in the real price of oil in August 1990, we support the view that the more restrictive Type A variant is adequate. However, changing from Type A and Type B can be a useful way of expressing different degrees of confidence in the narrative information itself.

B Robustness of Results for Monetary Policy Shocks

Note that the Narrative Sign Restriction 5 in the main text is of Type B. It postulates that the absolute value of the contribution of the monetary policy shock is "larger than the sum of the

draws that satisfy the baseline restrictions are generated. Out of these, 4,500 satisfy the narrative sign restriction. We approximate their weight in the importance step by using one thousand draws.

absolute value of the contribution of all other structural shocks" to the unexpected movement in the federal funds rate in October 1979. A Type A restriction would postulate that the contribution is "larger than the absolute value of the contribution of any other structural shocks." Clearly, in this case Type B is a stronger version than Type A. In our view, there is overwhelming evidence that the unexpected increase in the federal funds rate observed in October 1979 was the outcome of a monetary policy shock; hence, a Type B restriction is justified. Nevertheless, we will check the robustness of our results to specifying a milder Type A version of this restriction. To do this we will consider Alternative Narrative Sign Restriction 5.[5]

Alternative Narrative Sign Restriction 5. *For the observation corresponding to October 1979, a monetary policy shock is the most important driver of the unexpected movement in the federal funds rate. In other words, the absolute value of the contribution of monetary policy shocks to the unexpected movement in the federal funds rate is larger than the absolute value of the contribution of any other structural shock.*

Alternative Narrative Sign Restriction 5 does not meaningfully change the implications for the period around October 1979 relative to Narrative Sign Restriction 5. Figure B.1 replicates the panels displayed in Figure 5 and Figure B.2 replicates the IRFs displayed in Figure 6, but this time using Alternative Narrative Sign Restriction 5 instead of Narrative Sign Restriction 5. As the reader can see, the results are almost identical. Since Alternative Narrative Sign Restriction 5 is weaker than Narrative Sign Restriction 5, the contribution of the monetary policy shock is now slightly smaller and it is only responsible for between 50 and 115 basis points of the 225-basis-point unexpected increase in the federal funds rate observed in October of 1979.

[5]Narrative Sign Restriction 4 and Alternative Narrative Sign Restriction 5 affect in total one structural shock. The 10,116 draws generated in the previous exercise are used as the baseline. Out of these, 2,175 additionally satisfy Narrative Sign Restrictions 4 and Alternative Narrative Sign Restriction 5. We approximate their weights in the importance step by using one thousand draws.

Figure B.1: Results Around October 1979 with Narrative Sign Restrictions

(Narrative Sign Restriction 4 and Alternative Narrative Sign Restriction 5)

(a) Monetary Shock for October 1979

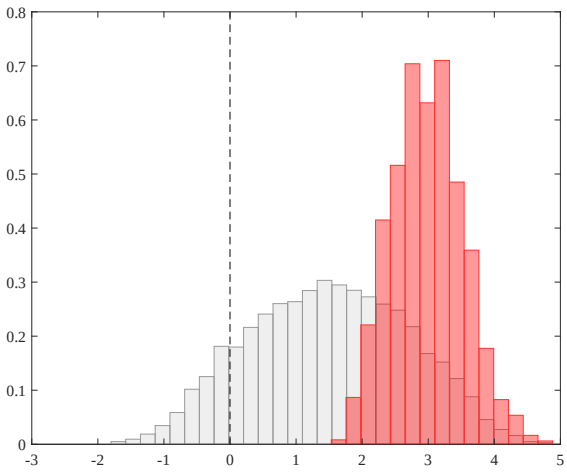

(b) Counterfactual Federal Funds Rate

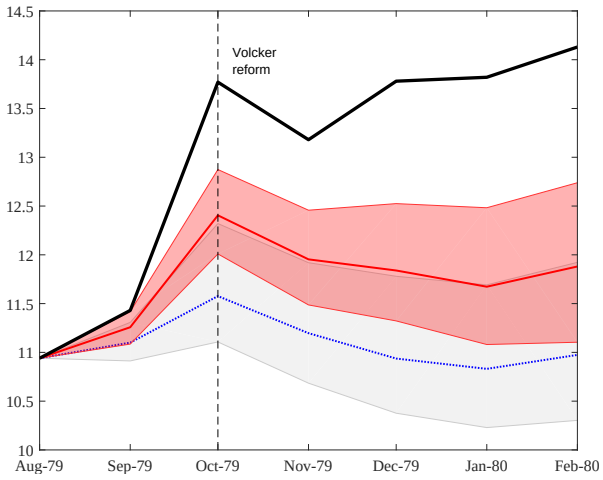

Note: Panel (a) plots the posterior distribution of the monetary policy shock for October 1979. Panel (b) plots the actual federal funds rate (solid wide) and the median of the counterfactual federal funds rate (solid thin) resulting from excluding all non-monetary structural shocks. The light bands represent 68% (point-wise) HPD credible sets around the median.

Figure B.2: IRFs with and without Narrative Sign Restrictions

(Narrative Sign Restriction 4 and Alternative Narrative Sign Restriction 5)

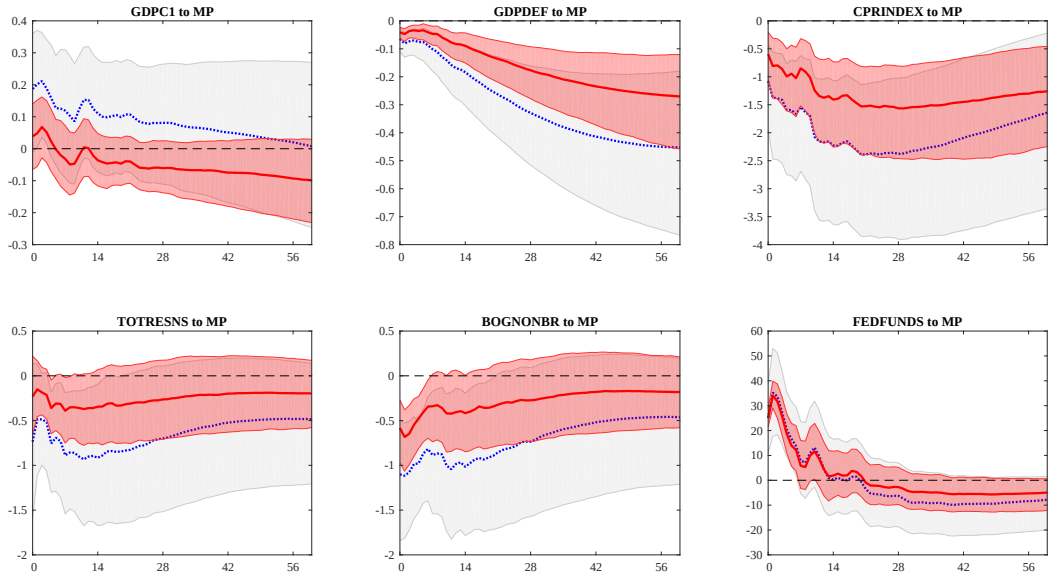

Note: The light shaded area represents the 68% (point-wise) HPD credible sets for the IRFs and the dotted lines are the median IRFs using the baseline identification restrictions. The darker shaded areas and solid lines display the equivalent quantities for the models that additionally satisfy Restriction Narrative Sign Restriction 4 and Alternative Narrative Sign Restriction 5. The IRFs have been normalized so that the monetary policy shock has an impact of 25 basis points on the federal funds rate.

C A New Chronology of Monetary Policy Shocks

In Section 6 we showed that using narrative information on a single event – October 1979 – is enough to obtain the result that the effect of contractionary monetary policy shocks on output are negative with very high posterior probability. That event is in our view the clearest and most uncontroversial, but there is a long literature that uses narrative and historical sources to isolate monetary policy shocks. This section first checks whether additional uncontroversial narrative information is available and second whether imposing it sharpens the results.

Following the pioneering work of Friedman and Schwartz (1963), Romer and Romer (1989) (henceforth, RR-89) combed through the minutes of the FOMC to create a dummy series of events which they argued represented exogenous tightenings of monetary policy. Focusing exclusively on contractionary shocks, they singled out a handful of episodes in the postwar period "in which the Federal Reserve attempted to exert a contractionary influence on the economy in order to reduce inflation" (RR-89, p. 134). These are December 1968, April 1974, August 1978 and October 1979. Romer and Romer (1994) later added December 1988 to the list. The Romer and Romer (1989, 1994) monetary policy narrative became very influential, but has been criticized by Leeper (1997), who pointed out that their dates are predictable from past macroeconomic data. As a consequence, in recent years alternative methods have been developed to construct time series of monetary policy shocks that are by design exogenous to the information set available at the time of the policy decision. The first prominent example is Romer and Romer (2004) (henceforth, RR-04), who regressed changes of the intended federal funds rate between FOMC meetings on changes in the Fed's Greenbook forecasts of output and inflation. By construction, the residuals from this regression are orthogonal to all the information contained in the Greenbook forecasts, and can plausibly be taken to be a measure of exogenous monetary policy shocks. A second approach looks at high-frequency financial data. Gürkaynak et al. (2005) look at movements in federal

funds futures contracts during a short window around the time of policy announcements to isolate the monetary policy shocks.

Figure C.1: Chronology of Monetary Policy Shocks

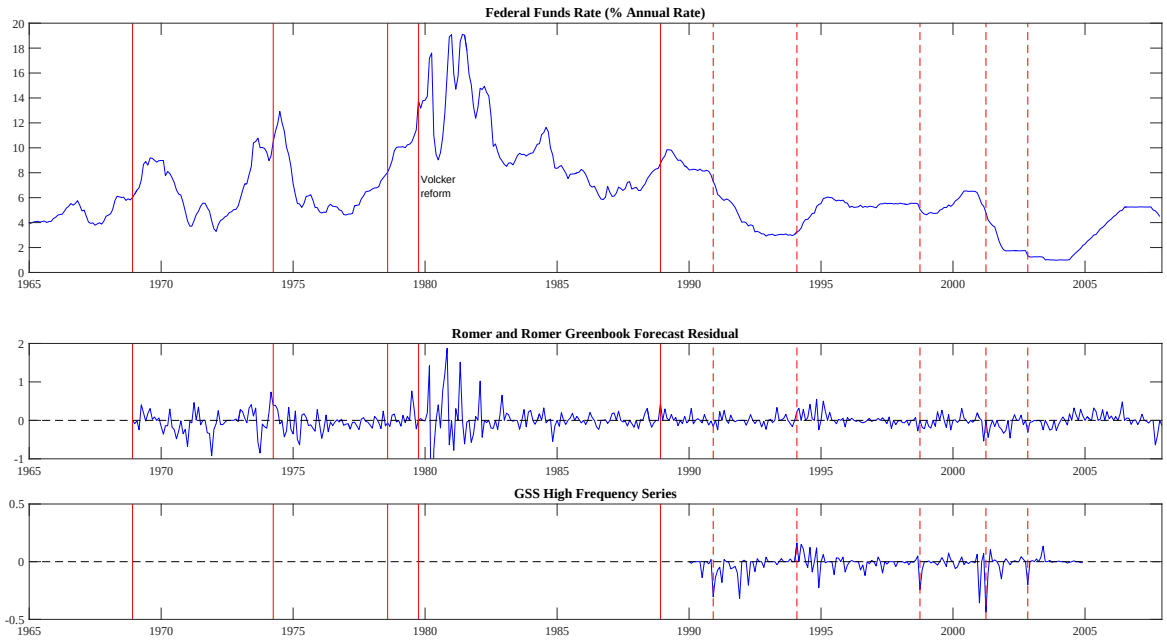

Note: The upper panel displays the average monthly level of the effective federal funds rate, in percent annual terms. The middle panel displays the Romer and Romer (2004) Greenbook forecast residual series, extended to 2007, while the lower panel displays the Gürkaynak et al. (2005) federal funds surprise series. The solid vertical lines represent the original dates singled out as monetary policy shocks by Romer and Romer, (1989, 1994), whereas the dashed vertical lines represent the additional episodes identified in the chronology below.

The solid vertical lines in Figure C.1 represent the Romer and Romer (1989, 1994) dates. The middle panel plots the RR-04 residuals, extended backward one month to cover the December 1968 meeting and forward to the end of 2007, whereas the lower panel plots the Gürkaynak et al. (2005) measure of monetary policy shocks. During the subsamples in which the series overlap, they disagree a great deal. For this reason, we will draw on the three approaches to select the dates for which the evidence of an exogenous monetary policy shock appears to be most compelling. For the period 1965-1991, which overlaps with the analysis conducted by Romer and Romer (1989, 1994), we start with their dates as candidate shocks

and review the evidence in light of the RR-04 analysis. Of the five Romer and Romer (1989, 1994) dates we keep three. The reasons for the choice are as follow.

- *December 1968.* After remaining stable around 6% for much of 1968, the federal funds rate began increasing gradually after the December meeting, a tightening that accelerated in the spring of 1969. It is unclear, however, that this event qualifies as a monetary policy shock. RR-89 (p. 140, footnote 13) recognize that "the tightening that occurred in December was in part a response to evidence of stronger growth," and the updated Greenbook residual series shows no shock for that meeting, suggesting that the roughly 25-basis-point increase in the federal funds rate registered that month can be fully explained by stronger output and inflation forecasts. We therefore exclude this event from our chronology.

- *April 1974.* Facing weak economic activity and accelerating inflation after the 1973 Arab oil embargo, the Fed chose to tighten policy, allowing the federal funds rate to rise to about 12% before loosening again with the objective of countering inflation expectations. The analysis of the Greenbook forecast reveals an outsized response of the Fed to the prevailing macroeconomic conditions. Indeed, the RR-04 series displays large positive residuals around this event, making it a good candidate for a monetary policy shock.

- *August 1978.* While RR-89 point to this event as an exogenous monetary policy tightening, an analysis of the Greenbook forecasts suggests that in fact much of this tightening can be explained by the Fed's systematic response to output and inflation. Indeed, the inflation outlook had deteriorated consistently in the spring and early summer of 1978, and the RR-04 series suggests that policy was broadly neutral, if not slightly loose, in August 1978 and subsequent months. We therefore exclude this event from our chronology.

- *October 1979.* The monetary policy decisions of October 6, 1979, enacted shortly after Paul Volcker became chairman of the Fed, are described by RR-89 as "a major anti-inflationary shock to monetary policy," and represent in our view the clearest case in the postwar period of an exogenous policy shock. Lindsey et al. (2013) provide a detailed narrative account of the events leading to the decision to abandon targeting the federal funds rate in favor of targeting nonborrowed reserves as the operating procedure for controlling the money supply. While macroeconomic conditions and, in particular, the deterioration of the inflation outlook and the increase in oil prices that followed the Iranian Revolution of 1978-79 played a large role in causing the shift, the forcefulness of the action, the surprise character of the action, and the dramatic break with established practice in the conduct of policy strongly suggest the occurrence of a monetary policy shock.[6]

- *December 1988.* Romer and Romer (1994) extended the original RR-89 chronology to include the sequence of interest rate increases that started in late 1988. As in previous events, their examination of the records of policy points to a shift toward tighter policy in order to "permit progress towards reducing inflation over time." This is confirmed by the Greenbook series, which shows that inflation forecasts did not worsen during that period, and real growth forecasts were revised upwards only moderately. Indeed the RR-04 series displays a positive value of 44 basis points in December 1988 and additional positive values for the subsequent four months. Therefore, the evidence appears to favor the occurrence of a monetary policy shock during this period.

We now turn to the 1990-2007 period, which was not covered by the Romers' original chronology. This period poses additional challenges given that, as argued by Ramey (2016),

[6]Note that because the RR-04 measure by construction includes only decisions that were made at regularly scheduled FOMC meetings, and the October 1979 reform was announced on a Saturday and outside of the regular FOMC cycle, the observation corresponding to this period is not available in the RR-04 series.

monetary policy has been conducted in a more systematic way, so true monetary policy shocks are now rare and therefore harder to identify. It is difficult to find instances that match the Romers' criterion of an event in which the Fed attempted to engineer a recession in order to bring down inflation, since inflation has been low and stable since the early 1990s. There are, however, a number of instances in which the Fed deviated from its usual behavior, responding more aggressively than normal in order to offset perceived risks to its inflation and employment goals. By construction, both the RR-04 measure and the high-frequency measure of Gürkaynak et al. (2005) (henceforth, GSS), which are available for this period, are likely to capture this type of event well. We single out as events December 1990, February 1994, October 1998, January 2001 and November 2002. These are represented by the dashed vertical lines in Figure C.1. With the exception of the 1994 event, they all represent circumstances in which the Fed eased aggressively, citing "risk management" considerations in response to unusual risks to economic growth. Here we explain why we choose each of these five events.

- *December 1990.* During the fall of 1990 the FOMC had started to ease monetary policy in response to the Gulf War and the associated spike in oil prices, which was expected to cause an economic contraction. By the time of the FOMC meeting of December 18, hopes of a quick resolution of the war emerged and oil prices had reversed almost half of their increase. The Greenbook forecasts presented by the staff foresaw a more robust recovery during the subsequent spring, and the forecast for the level of output was revised upward for both the December and the February meetings. The FOMC, however, decided to ease policy further on both occasions, contrary to expectations (as seen by the presence of negative shocks in the GSS series) and to its usual reaction function (as seen in the RR-04 series), citing the need to "insure" the economy from the risk of a deeper recession or further shocks.[7] We will therefore keep this event in

[7]The main justification for the surprisingly dovish stance appears to be an unwillingness to sacrifice output in

the chronology.

- *February 1994.* Starting in February 1994, the FOMC began a series of tightening moves that over the subsequent 12 months increased the fed funds rate by 300 basis points. The start of the tightening campaign certainly was a complete surprise to financial market participants, leading to a major adjustment in longer-term interest rates.[8] The speed of subsequent hikes was also a surprise, as can be seen from the GSS series. Moreover, the sequence of interest rate increases appears aggressive relative to usual procedures. Indeed, the RR-04 series displays a positive shock for the observation corresponding to every single meeting up to November 1994, and an examination of the staff projections and forecasts prepared for the Feburary meeting reveals that the tightening between February and November was more aggressive than both the baseline policy proposal prepared by the staff, and a tighter policy alternative. There is evidence, however, that the 1994 event could be an example of superior information, or "policy foresight," rather than a true monetary policy shock. Indeed, an examination of the minutes of the February 1994 FOMC meeting reveals that policy makers had confidential access to the employment data to be released publicly later that day, and which had not been available for the preparation of the Greenbook forecast, indicating that at least part of the tightening was a response to news on improving economic activity. Nevertheless, the minutes of the FOMC meetings in the early part of 1994 do reveal an outsized response to the risk of inflation accelerating.

- *October 1998.* In late September of 1998, the FOMC responded to the deterioration in

order to reduce inflation. "While substantial additional easing might not be needed under prevailing conditions, a limited further move would provide some added insurance in cushioning the economy against the possibility of a deepening recession and an inadequate rebound in the economy without imposing an unwarranted risk of stimulating inflation later."

[8]See "The great bond massacre" (Fortune, 1994) for a representative contemporary account, which associated the heavy losses experienced by financial companies, hedge funds, and bond mutual funds on their holdings of long-term bonds with the surprise tightening by the Fed.

the global economic outlook stemming from the Russian debt crisis of 1998 and the failure of the hedge fund Long Term Capital Management (LTCM) by lowering the federal funds rate by 25 basis points "to cushion the effects on prospective economic growth in the United States of increasing weakness in foreign economies and of less accommodative financial conditions domestically."[9] On October 15, after an unscheduled intermeeting conference call two weeks later, the FOMC decided to cut by an additional 25 basis points. As can be seen from the GSS series, the move came as a surprise to financial markets. An examination of the transcript of the conference call reveals that there had not been material changes to economic data in the prior two weeks, and that the FOMC was deliberating on "a matter of uncertainties at this point [rather] than clear-cut changes in the outlook," on the basis of turbulence in financial markets. A participant in the meeting pointed out that there was "no basis there for a material change in policy," but "a higher degree of uncertainty [which] reinforces the sense of downside risks."[10] This episode in which the FOMC was seen to respond to financial turbulence alone led to the expression "Greenspan 'put'," which referred to the perceived insurance the Fed was providing to financial market participants against stock market crashes. We will therefore keep this event in the chronology.

- *April 2001.* In response to the weakening in the economy that had begun in the fall of 2000, the Federal Reserve began lowering the federal funds rate with a 50-basis-point cut on January 3, 2001. While the timing of the move was a surprise (it took place during an intermeeting conference call shortly after taking no action at the December meeting just a few weeks earlier), it is unclear whether the January cut can be classified as a monetary policy shock. All participants in the meeting explicitly mentioned a deteriorating outlook for the economy as the reason for lowering interest rates.

[9] See *Statement*, Federal Open Market Committee, September 29, 1998.
[10] See *Transcript*, Federal Open Market Committee, October 15, 1998.

Moreover, in the transcript of the conference call, Chairman Greenspan explicitly mentions having received classified data on unemployment claims pointing to further weakness. A stronger case can be built for the April 18, 2001 meeting, another instance of the FOMC lowering the federal funds rate in a surprise move in between scheduled meetings. In his opening statement, Chairman Greenspan made clear that "in reviewing the economic outlook over the last week, itâĂŹs fairly apparent that very little of significance has changed." It appears that during this period, as in the 1998 episode, the FOMC was placing a substantial weight on asset price volatility, particularly after the bursting of the dot-com stock price bubble the previous year. We will therefore keep this event in the chronology.

- *November 2002.* In November of 2002 the FOMC lowered the federal funds rate by 50 basis points. This move was both larger than what the market expected, and what, according to the updated RR-04 Greenbook series, was warranted by the available economic data. Moreover, incoming data received after the completion of the Greenbook "were very close to our expectations and require little change to [the] near-term forecast." Particularly in light of developments in Japan, which had been experiencing persistent deflation since the late 1990s, it appears that concerns about deflation loomed large.[11] Geopolitical risks – preparations for the 2003 Iraq war were already under way – were also a concern.[12] Once again, risk management considerations motivated a larger-than-usual cut that would provide 'insurance against downside risks." As Chairman Greenspan argued, "If we move significantly today –and my suggestion would be to lower the funds rate 50 basis pointsâĂŤ- it is possible that

[11]One participant expressed concern that "a negative demand shock could cause the disinflation trends we've had lately to morph into deflation," and staff simulations placed a 25-30% probability that the economy would experience a deflation. Chairman Greenspan remarked that "if we were to fail to move and the economy began to deteriorate [...] we were looking into a deep deflationary hole." See *Transcript*, Federal Open Market Committee, November 6, 2002.

[12]See *Statement*, Federal Open Market Committee, November 6, 2002.

such a move may be a mistake. But it's a mistake that does not have very significant consequences. On the other hand, if we fail to move and we are wrong, meaning that we needed to, the cost could be quite high."[13] We will therefore keep this event in the chronology.

To summarize, by cross-checking the updated Greenbook residual series from RR-04, the high-frequency series from GSS, and the transcripts from the meetings of the FOMC, we have singled out eight events for which there appears to be a good case that a monetary policy shock occurred. Of these, four were contractionary, or tightening, shocks (positive in terms of their impact on the federal funds rate) and four were expansionary, or easing, shocks (negative shocks). We will therefore consider the Narrative Sign Restriction 6 and 7 stated in Section 6, reproduced below for convenience:

Narrative Sign Restriction 6. *The monetary policy shock must be positive for the observations corresponding to April 1974, October 1979, December 1988, and February 1994, and negative for December 1990, October 1998, April 2001, and November 2002.*

Narrative Sign Restriction 7. *For the periods specified by Restriction 6, monetary policy shocks are the most important contributor to the observed unexpected movements in the federal funds rate. In other words, the absolute value of the contribution of monetary policy shocks is larger than the absolute value of the contribution of any other structural shock.*

In terms of the definitions of Section 3, Narrative Sign Restriction 6 is a restriction on the sign of the structural shocks, whereas Narrative Sign Restriction 7 is a Type A restriction on the historical decomposition.

Figure C.2 presents the IRFs to a monetary policy shock, with and without narrative information. The light shaded area represents the 68% (point-wise) HPD credible sets for the

[13] See *Transcript*, Federal Open Market Committee, November 6, 2002.

Figure C.2: IRFs with and without Narrative Sign Restrictions

(Narrative Sign Restrictions 6, and 7)

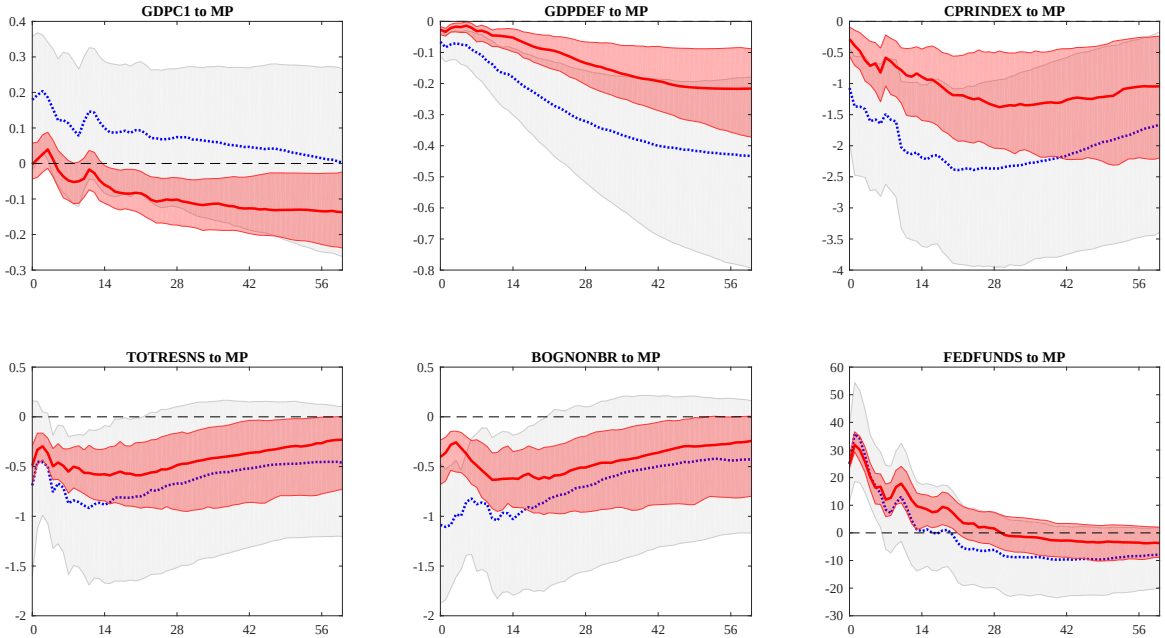

Note: The light shaded area represents the 68% (point-wise) HPD credible sets for the IRFs and the dotted lines are the median IRFs using the baseline identification restrictions. The darker shaded areas and solid lines display the equivalent quantities for the models that additionally satisfy Narrative Sign Restrictions 6-7.

IRFs and the dotted lines are the median IRFs using the baseline identification. These results replicate the IRFs depicted in Figure 6 of Uhlig (2005). The darker shaded areas and solid lines display the equivalent quantities when Narrative Sign Restrictions 6-7 are also used. The results are very similar to those using only the Volcker episode, reported in Figure 6 in the main text.

C.1 Results for the 1994 event only

As mentioned in the main text, the February 1994 it is worth highlighting that the February 1994 event stands out because the narrative record identifies a major contractionary monetary shock, which was followed in the data by a subsequent boom in output. Therefore, the historical account arguably cannot be considered to be distorted by the presence of a

recession shortly after the monetary policy action. We can specify the following narrative sign restrictions:

Narrative Sign Restriction 8. *The monetary policy shock must be positive for the observation corresponding to February 1994.*

Narrative Sign Restriction 9. *For the period specified by Restriction 8, the monetary policy shock is the overwhelming contributor to the observed unexpected movements in the federal funds rate. In other words, the absolute value of the contribution of monetary policy shocks is larger than the sum of the absolute value of the contribution of all other structural shock.*

In terms of the definitions of Section 3, Narrative Sign Restriction 8 is a restriction on the sign of the structural shocks, whereas Narrative Sign Restriction 9 is a Type A restriction on the historical decomposition. Figure C.3 displays the results. As can be seen, the results are qualitatively similar to the ones using the Volcker event. The response of GDP is negative, and significantly after about a year. Thus, we conclude that it is possible to find an event that did not lead to a recession shortly thereafter – in fact, output growth accelerated in 1994 and 1995 – and yet contains useful information for identification.

Figure C.3: IRFs with and without Narrative Sign Restrictions

(1994 event only)

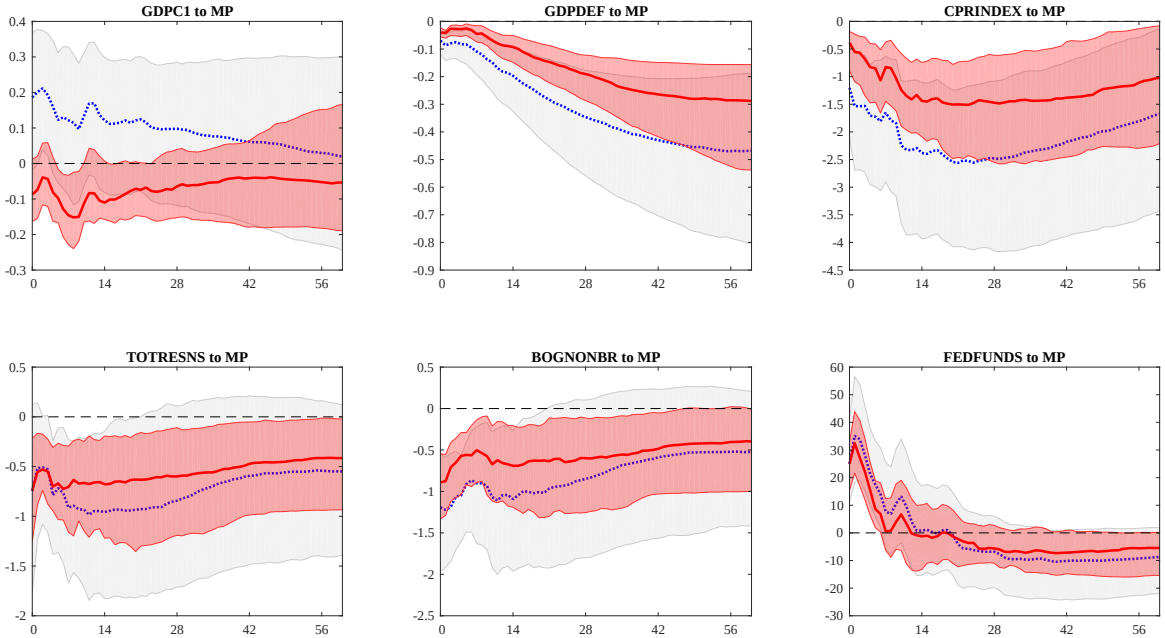

Note: The light shaded area represents the 68% (point-wise) HPD credible sets for the IRFs and the dotted lines are the median IRFs using the baseline identification restrictions. The darker shaded areas and solid lines display the equivalent quantities for the models that additionally satisfy Narrative Sign Restrictions 6-7.

D Computational Aspects

In this section we give additional details on the computational properties of our Algorithm 1 for the applications presented in the text. First, recall that Steps 1, 2, and 4 are basically identical to the steps in standard algorithms with traditional sign restrictions; see Rubio-Ramirez et al. (2010) and Arias et al. (2016b). The only difference is that now we also need to check whether the narrative sign restrictions hold in addition to checking whether the traditional sign restrictions hold. Second, the main difference with standard algorithms with only traditional sign restrictions is Step 3. After discarding draws that do not satisfy the restrictions, in Step 3 we compute the weights for the accepted draws and in Step 5 we draw with replacement using those weights. Third, Step 1 produces independent draws. This is

very advantageous as the number of draws needed to conduct inference will be much lower than with less efficient algorithms, for instance, those based on Metropolis-Hastings. Fourth, the rejection rate in Step 2 will vary from application to application. We have calibrated the number of draws in Step 1 to obtain around $1,000$ draws that satisfy both the traditional and narrative sign restrictions. This implies that when imposing one or two narrative sign restrictions as in the results presented in Sections 5.4 or 6.3, we only need to produce $10,000$ draws in Step 1. When more narrative sign restrictions are used, as in the results presented in Section 5.3, we might need to obtain more than 10^8 draws in Step 1. Finally, the reader should notice that given that the algorithm is easy to parallelize, obtaining 10^8 draws in Step 1 is not a problem. The computations in this paper were carried using MATLAB R2016b on an 12-core HP Z420 computer with an Intel Xeon CPU with a 360Ghz processor and 16 Gb of RAM and 10^8 draws were obtained in less than one hour.

Having said this, because we use an importance sampling step, we need to compute the effective sample size to guard against having only a few draws dominate Step 5. If w_i is the weight associated with the ith draw, then the effective sample size is

$$\text{ESS} = \frac{\left(\sum_{i=1}^{N} w_i\right)^2}{\sum_{i=1}^{N} w_i^2}$$

where N is the number of draws that satisfy the restrictions. The effective sample size can be expressed as a percentage by dividing ESS by N and multiplying by 100. For all the applications in the paper we obtained ESS (expressed in percentage) of 85% or above. Hence we do not see any issues with the importance sampling weights, even in the case of considering several narrative sign restrictions as it is in the case of Section 5.3.

Finally, it is important to mention that the number of draws M necessary to compute weights also depends on the number of narrative sign restrictions considered. In the case of many narrative sign restrictions, as Section 5.3, we need up to 10^6 draws. In the case of

only one narrative sign restriction as in the results presented in Sections 5.4 or 6.3, $1,000$ draws are enough. Computing the weights using 10^6 draws maybe time consuming, but the reader should remember that we only need to compute the weights for the draws that satisfy both the traditional and narrative sign restrictions, i.e. around $1,000$ times. When 10^6 draws were needed to compute the $1,000$ weights, our code needed 3.5 hours to do so. When $1,000$ draws were needed to compute the $1,000$ weights, our code took less than one minute. Once again, since the draws of the algorithm are independent, these times can be massively accelerated using modern parallel computing techniques.