

Antolín-Díaz, Juan; Rubio-Ramírez, Juan Francisco

Working Paper

Narrative sign restrictions for SVARs

Working Paper, No. 2016-16

Provided in Cooperation with:

Federal Reserve Bank of Atlanta

Suggested Citation: Antolín-Díaz, Juan; Rubio-Ramírez, Juan Francisco (2016) : Narrative sign restrictions for SVARs, Working Paper, No. 2016-16, Federal Reserve Bank of Atlanta, Atlanta, GA

This Version is available at:

<https://hdl.handle.net/10419/172913>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Narrative Sign Restrictions for SVARs

Juan Antolín-Díaz and Juan F. Rubio-Ramírez

Working Paper 2016-16

December 2016

Abstract: We identify structural vector autoregressions using narrative sign restrictions. Narrative sign restrictions constrain the structural shocks and 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 are highly informative. We highlight that adding a single narrative sign restriction dramatically sharpens and even changes 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 Waggoner, 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.

Federal Reserve Bank of Atlanta working papers, including revised versions, are available on the Atlanta Fed's website at www.frbatlanta.org. Click "Publications" and then "Working Papers." To receive e-mail notifications about new papers, use frbatlanta.org/forms/subscribe.

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 few additional uncontentious sign restrictions that help shrink the set of admissible structural parameters and allow us to reach clear economic conclusions.

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 historical decomposition agree with the established narrative. For example, narrative sign restrictions will 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” or 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, the narrative sign restrictions instead truncate the support of the likelihood function. Crucially, the Bayesian methods developed in Rubio-Ramirez et al. (2010) and Arias et al. (2016b) only need to be slightly modified for the case of narrative sign restrictions. We see narrative sign restrictions as complementing the traditional ones, and in our empirical applications we will combine both.

There is a long tradition, starting with Friedman and Schwartz (1963), of using 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 and Shapiro, 1998, 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 “true” 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 in the model. 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

¹See also Montiel-Olea et al. (2015), Caldara and Herbst (2016), and Drautzburg (2016).

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. This alleviates the issue of measurement error in the narrative time series, since 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, adding narrative sign restrictions based on a small set of historical events allows us to distinguish between aggregate demand and oil-specific demand shocks. In fact, we show that 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) 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.

The rest of this paper is organized as follows. Section 2 presents the basic SVAR framework. Section 3 lays out the identification problem and shows how to incorporate narrative information as sign restrictions. Section 4 describes the algorithms. Section 5 applies the methodology to the oil market, while Section 6 applies the methodology to the identification of monetary policy shocks.

Section 7 concludes.

2 The Model

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

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

where \mathbf{y}_t is an $n \times 1$ vector of variables, $\boldsymbol{\varepsilon}_t$ is an $n \times 1$ vector of structural shocks, \mathbf{A}_ℓ 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 $\boldsymbol{\varepsilon}_t$, conditional on past information and the initial conditions $\mathbf{y}_0, \dots, \mathbf{y}_{1-p}$, 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

$$\mathbf{y}'_t \mathbf{A}_0 = \mathbf{x}'_t \mathbf{A}_+ + \boldsymbol{\varepsilon}'_t \quad \text{for } 1 \leq t \leq T, \quad (2)$$

where $\mathbf{A}'_+ = [\mathbf{A}'_1 \cdots \mathbf{A}'_p \ \mathbf{c}']$ and $\mathbf{x}'_t = [\mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-p}, 1]$ 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

$$\mathbf{y}'_t = \mathbf{x}'_t \mathbf{B} + \mathbf{u}'_t \quad \text{for } 1 \leq t \leq T,$$

where $\mathbf{B} = \mathbf{A}_+ \mathbf{A}_0^{-1}$, $\mathbf{u}'_t = \boldsymbol{\varepsilon}'_t \mathbf{A}_0^{-1}$, and $\mathbb{E}[\mathbf{u}_t \mathbf{u}'_t] = \boldsymbol{\Sigma} = (\mathbf{A}_0 \mathbf{A}'_0)^{-1}$. The matrices \mathbf{B} and $\boldsymbol{\Sigma}$ are the reduced-form parameters, while \mathbf{A}_0 and \mathbf{A}_+ are the structural parameters. Let $\boldsymbol{\Theta} = (\mathbf{A}_0, \mathbf{A}_+)$ collect the value of the structural parameters.

2.1 Impulse response functions

Recall the definition of impulse response functions (IRFs).

Definition 1. Given a value Θ of the structural parameters, the IRF 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 $\mathbf{L}_k(\Theta)$, where $\mathbf{L}_k(\Theta)$ is defined recursively by

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

We will refer to the entry (i, j) of the matrix $\mathbf{L}_k(\Theta)$ as $\mathbf{L}_k(\Theta)_{i,j}$.

2.2 Structural shocks and historical decomposition

Given a value Θ of the structural parameters and the data, the structural shocks at time t are defined by

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

The historical decomposition calculates the cumulative contribution of the structural shocks to the observed unexpected change in the variables between two periods. Formally, the contribution of the j -th structural shock to the observed unexpected change in the i -th variable between periods t and $t+h$ is

$$H_{i,j,t,t+h}(\Theta, \boldsymbol{\varepsilon}_t, \dots, \boldsymbol{\varepsilon}_{t+h}) = \sum_{\ell=0}^h \mathbf{e}'_{i,n} \mathbf{L}_\ell(\Theta) \mathbf{e}_{j,n} \mathbf{e}'_{j,n} \boldsymbol{\varepsilon}_{t+h-\ell},$$

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. Sign restrictions will lead to the structural parameters being set-identified, not point-identified. The set of admissible structural parameters will often be very large and its elements will have very different implications for IRFs, elasticities, historical decompositions or forecasting error variance decompositions. Thus, being minimalist frequently comes at the cost of not arriving at meaningful economic conclusions. This point was first illustrated by Kilian and Murphy (2012), who used an SVAR of the oil market identified only through a few well-accepted sign restrictions on IRFs to show that the set of admissible structural parameters retains many with implausible implications for the price elasticity of oil supply. More recently, Arias et al. (2016a) have pointed out that the identification scheme of Uhlig (2005) retains many structural parameters with unconvincing implications for the systematic response of monetary policy to output. Hence, the challenge is to come up with additional uncontroversial sign restrictions that rule out some of these structural parameters and help sharpen inference.

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 historical decomposition 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 use narrative sign restrictions 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 aggregate demand shock) to the unexpected drop in oil production observed during that period.

We now 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) and highlight how the second depend on the structural shocks.

3.1 Sign restrictions on the impulse response functions

Let us assume that we want to identify the j -th structural shock by imposing s_j sign restrictions on the IRFs at different horizons. Define $\mathbf{F}(\Theta)$ as vertically stacking the IRFs at the different horizons over which we want to impose the restrictions and \mathbf{S}_j as an $s_j \times r$ matrix of zeros and ones that will select the horizons and the variables over which we want to impose the r sign restrictions. Then the values of the structural parameters that satisfy the sign restrictions are defined by

$$\phi_j(\Theta) = \mathbf{S}_j \mathbf{F}(\Theta) \mathbf{e}_{j,n} > \mathbf{0} \text{ for } 1 \leq j \leq n. \quad (4)$$

As an example, if we choose $\mathbf{F}(\Theta) = \mathbf{L}_0(\Theta)$ and $\mathbf{S}_j = \mathbf{e}'_{2,n}$, we are imposing the sign restriction that the IRF at horizon zero of the second variable to the j -th structural shock is positive. If the choices are instead $\mathbf{F}(\Theta) = \mathbf{L}_0(\Theta)$ and $\mathbf{S}_j = -\mathbf{e}'_{2,n}$, we are imposing the sign restriction that the IRF at horizon zero of the second variable to the j -th structural shock is negative. If we define $\mathbf{F}(\Theta) = (\mathbf{L}_0(\Theta)' \ \mathbf{L}_1(\Theta)')'$ and $\mathbf{F}(\Theta)$ with $\mathbf{S}_j = (\mathbf{e}_{3,2n} \ \mathbf{e}_{2n,2n})'$ instead, we are imposing the sign restriction that the IRF at horizon zero of the third variable to the j -th structural shock is positive and the sign restriction that the IRF at horizon one of the n -th variable to the j -th structural shock is positive.

3.2 Sign restrictions on the structural parameters

If we want to identify the j -th structural shock by imposing s_j sign restrictions on the structural parameters, we can then define $\mathbf{F}(\Theta) = \Theta$ and \mathbf{S}_j as an $s_j \times r$ matrix of zeros and ones that will select entries of Θ over which we want to impose the sign restrictions. In this case Equation (4) will characterize the values of the structural parameters that satisfy the sign restrictions. As an example, if we choose $\mathbf{F}(\Theta) = \Theta$ and $\mathbf{S}_j = \mathbf{e}'_{2,n+m}$, we are imposing the sign restriction that the entry $(2, j)$ of \mathbf{A}_0 is positive. If we choose $\mathbf{F}(\Theta) = \Theta$ with $\mathbf{S}_j = (\mathbf{e}_{\tilde{n},n+m} \mathbf{e}_{\hat{n},n+m})'$ instead, we are imposing the sign restrictions that the entries (\tilde{n}, j) and (\hat{n}, j) of Θ are positive, with $1 \leq \tilde{n}, \hat{n} \leq n + m$. As before, introducing negative sign restrictions simply requires using -1 instead of 1 in \mathbf{S}_j .

We now turn to narrative sign restrictions. We will first describe the function that characterizes narrative sign restrictions based on the signs of the structural shocks. Second, we will do the same for narrative sign restrictions based on the historical decomposition.

3.3 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 identify the j -th structural shock by imposing the restriction that the signs of the j -th structural shock at s_j episodes occurring at dates t_1, \dots, t_{s_j} are positive. Then, the narrative sign restrictions can be imposed as

$$\mathbf{e}'_{j,n} \boldsymbol{\varepsilon}_{t_v}(\Theta) > 0 \text{ for } 1 \leq v \leq s_j. \quad (5)$$

Assume instead that we want to identify the j -th structural shock by imposing the restriction that the signs of the j -th structural shock at s_j episodes occurring at dates t_1, \dots, t_{s_j} are negative. Then, the narrative sign restrictions can be imposed as

$$\mathbf{e}'_{j,n} \boldsymbol{\varepsilon}_{t_v}(\Theta) < 0 \text{ for } 1 \leq v \leq s_j. \quad (6)$$

Of course, Equations (5) and (6) can be used jointly. One could restrict a few structural shocks to

be negative and a few others to be positive.

3.4 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 structural shock was the most important shock for the unexpected movement of some variable during those periods. In particular, this is information on the relative magnitude of the contribution of the j -th structural shock to the unexpected change in the i -th variable between some periods. We propose to formalize this idea in two different ways. First, we may specify that a given structural 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 structural 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.

3.4.1 Type A restrictions on the historical decomposition

To fix ideas, let us consider that we want to identify the j -th structural shock by imposing the restriction that the absolute value of the contribution of the j -th structural shock to the unexpected change in the i -th variable between periods t and $t + h$ is larger than the absolute value of the contribution of any other structural shock to the unexpected change in the i -th variable between periods t and $t + h$. Then, the narrative sign restrictions can be imposed as

$$|H_{i,j,t,t+h}(\Theta, \varepsilon_t(\Theta), \dots, \varepsilon_{t+h}(\Theta))| - \max_{j' \neq j} |H_{i,j',t,t+h}(\Theta, \varepsilon_t(\Theta), \dots, \varepsilon_{t+h}(\Theta))| > 0,$$

where $|H_{i,j,t,t+h}(\Theta, \varepsilon_t(\Theta), \dots, \varepsilon_{t+h}(\Theta))|$ is the absolute value of $H_{i,j,t,t+h}(\Theta, \varepsilon_t(\Theta), \dots, \varepsilon_{t+h}(\Theta))$.

For example, assume we have a model with three variables and we want to identify the 2nd structural shock by imposing the condition that between periods 6 and 7 the absolute value of the contribution of the 2nd structural shock to the unexpected change in the 3rd variable is larger than the contribution of any other structural shock. Then, the narrative sign restrictions can be imposed as

$$|H_{3,2,6,7}(\Theta, \varepsilon_6(\Theta), \varepsilon_7(\Theta))| - \max_{j' \neq 2} |H_{3,j',6,7}(\Theta, \varepsilon_6(\Theta), \varepsilon_7(\Theta))| > 0.$$

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

$$|H_{i_v,j,t_v,t_v+h_v}(\Theta, \varepsilon_{t_v}(\Theta), \dots, \varepsilon_{t_v+h_v}(\Theta))| - \max_{j' \neq j} |H_{i_v,j',t_v,t_v+h_v}(\Theta, \varepsilon_{t_v}(\Theta), \dots, \varepsilon_{t_v+h_v}(\Theta))| > 0, \quad (7)$$

for $1 \leq v \leq s_j$.

Assume instead that one wants to identify the j -th structural shock by imposing the restriction that the absolute value of the contribution of the j -th structural shock to the unexpected change in the i_1, \dots, i_{s_j} -th variables from periods t_1, \dots, t_{s_j} to $t_1 + h_1, \dots, t_{s_j} + h_{s_j}$ is smaller in absolute value than the contribution of any other structural shock to the unexpected change in those variables during those periods. Then, the narrative sign restrictions can be imposed as

$$|H_{i_v,j,t_v,t_v+h_v}(\Theta, \varepsilon_{t_v}(\Theta), \dots, \varepsilon_{t_v+h_v}(\Theta))| - \min_{j' \neq j} |H_{i_v,j',t_v,t_v+h_v}(\Theta, \varepsilon_{t_v}(\Theta), \dots, \varepsilon_{t_v+h_v}(\Theta))| < 0, \quad (8)$$

for $1 \leq v \leq s_j$. As mentioned above, Equations (7) and (8) can be used jointly.

3.4.2 Type B restrictions on the historical decomposition

As before, to fix ideas, assume we want to identify the j -th structural shock by imposing the restriction that the absolute value of the contribution of the j -th structural shock to the unexpected change in the i -th variable between periods t and $t + h$ is larger than the sum of the absolute value of the contribution of all other structural shocks to the unexpected change in the i -th variable between periods t and $t + h$. Then, the narrative sign restrictions can be imposed as

$$|H_{i,j,t,t+h}(\Theta, \varepsilon_t(\Theta), \dots, \varepsilon_{t+h}(\Theta))| - \sum_{j' \neq j} |H_{i,j',t,t+h}(\Theta, \varepsilon_t(\Theta), \dots, \varepsilon_{t+h}(\Theta))| > 0.$$

For example, assume we have a model with three variables and we want to identify the 2nd structural shock by imposing the restriction that the absolute value of the contribution of the 2nd structural shock to the unexpected change in the 3rd variable between periods 6 and 7 is larger than the sum of the absolute value of the contribution of all other structural shocks to the unexpected change in the 3rd variable between periods 6 and 7. Then, the narrative sign restrictions can be imposed as

$$|H_{3,2,6,7}(\Theta, \varepsilon_6(\Theta), \varepsilon_7(\Theta))| - \sum_{j' \neq 2} |H_{3,j',6,7}(\Theta, \varepsilon_6(\Theta), \varepsilon_7(\Theta))| > 0.$$

As before, we can identify the j -th structural shock by imposing s_j restrictions of this type. Suppose that we want to identify the j -th structural shock by imposing the restriction that the absolute value of the contribution of the j -th structural shock to the unexpected change in the i_1, \dots, i_{s_j} -th variables from periods t_1, \dots, t_{s_j} to $t_1 + h_1, \dots, t_{s_j} + h_{s_j}$ is larger than the sum of the absolute values of the contributions of all other structural shocks to the unexpected change in those variables and for those periods. Then, we can define

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

for $1 \leq v \leq s_j$.

Suppose now that we want to identify the j -th structural shock by imposing the restriction that the absolute value of the contribution of the j -th structural shock to the unexpected change in the i_1, \dots, i_{s_j} -th variables from periods t_1, \dots, t_{s_j} to $t_1 + h_1, \dots, t_{s_j} + h_{s_j}$ is smaller than the sum of the absolute values of the contributions of all other structural shocks to the unexpected change in those variables and for those periods. Then, we can define

$$|H_{i_v, j, t_v, t_v+h_v}(\Theta, \varepsilon_{t_v}(\Theta), \dots, \varepsilon_{t_v+h_v}(\Theta))| - \sum_{j' \neq j} |H_{i_v, j', t_v, t_v+h_v}(\Theta, \varepsilon_{t_v}(\Theta), \dots, \varepsilon_{t_v+h_v}(\Theta))| < 0, \quad (10)$$

for $1 \leq v \leq s_j$. Equations (9) and (10) can also be used jointly.

3.4.3 Discussion

A natural question is to ask whether Type A or Type B restrictions on the historical decomposition of the data into structural shocks 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. On the contrary, 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. Whether Type A or Type B is more suitable needs to be decided on a case-by-case basis, depending on the level of confidence the researcher has in the narrative information about a particular episode.

4 Implementation

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)-(10) imply the

following function to characterize narrative sign restrictions

$$\phi(\Theta, \varepsilon^v) > \mathbf{0}, \quad (11)$$

where $\varepsilon^v = (\varepsilon_{t_1}, \dots, \varepsilon_{t_v})$ are the structural shocks constrained by the narrative sign restriction. Moreover, Equation (3) implies the following function

$$\varepsilon_t = g_h(\mathbf{y}_t, \mathbf{x}_t, \Theta) \text{ for } 1 \leq t \leq T, \quad (12)$$

which is invertible

$$\mathbf{y}_t = g_h^{-1}(\varepsilon_t; \mathbf{x}_t, \Theta) \text{ for } 1 \leq t \leq T.$$

Using Equations (11) and (12), we can write

$$\tilde{\phi}(\Theta, \mathbf{y}^v, \mathbf{x}^v) = \phi(\Theta, g_h(\mathbf{y}_{t_1}, \mathbf{x}_{t_1}, \Theta), \dots, g_h(\mathbf{y}_{t_v}, \mathbf{x}_{t_v}, \Theta)) > \mathbf{0}, \quad (13)$$

where $\mathbf{y}^v = (\mathbf{y}_{t_1}, \dots, \mathbf{y}_{t_v})$ and $\mathbf{x}^v = (\mathbf{x}_{t_1}, \dots, \mathbf{x}_{t_v})$. Hence, given the data, Equation (11) is continuous on the structural parameters while, given the structural parameters, Equation (11) is continuous on the structural shocks.

4.1 The posterior distribution

Following Arias et al. (2016b), we can consider an alternative parameterization of Equation (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 Θ 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 function f_h by

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

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

$$f_h^{-1}(\mathbf{B}, \Sigma, \mathbf{Q}) = \underbrace{(h(\Sigma)^{-1} \mathbf{Q})}_{\mathbf{A}_0}, \underbrace{\mathbf{B} h(\Sigma)^{-1} \mathbf{Q}}_{\mathbf{A}_+}. \quad (14)$$

Using Equation (14), we can rewrite Equation (13) as

$$\Phi(\mathbf{B}, \Sigma, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) = \tilde{\phi}(f_h^{-1}(\mathbf{B}, \Sigma, \mathbf{Q}), \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}.$$

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

$$\begin{aligned} & \pi(\mathbf{B}, \Sigma, \mathbf{Q} | \mathbf{y}^T, \Phi(\mathbf{B}, \Sigma, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}) \\ &= \frac{\pi(\mathbf{y}^T | \mathbf{B}, \Sigma, \mathbf{Q}, \Phi(\mathbf{B}, \Sigma, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}) \pi(\mathbf{B}, \Sigma, \mathbf{Q})}{\int \pi(\mathbf{y}^T | \mathbf{B}, \Sigma, \mathbf{Q}, \Phi(\mathbf{B}, \Sigma, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}) \pi(\mathbf{B}, \Sigma, \mathbf{Q}) d(\mathbf{B}, \Sigma, \mathbf{Q})} \end{aligned} \quad (15)$$

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

It is useful at this point to compare the posterior distribution defined in Equation (15) with the one obtained using only traditional sign restrictions. As defined by Equation (4), the latter are characterized by $\phi(\Theta) > \mathbf{0}$, which does not depend on the realization of the shock. Thus, the posterior of $(\mathbf{B}, \Sigma, \mathbf{Q})$ subject to the traditional sign restrictions is

$$\begin{aligned} & \pi(\mathbf{B}, \Sigma, \mathbf{Q} | \mathbf{y}^T, \phi(f_h^{-1}(\mathbf{B}, \Sigma, \mathbf{Q})) > \mathbf{0}) = \\ &= \frac{\pi(\mathbf{y}^T | \mathbf{B}, \Sigma) \pi(\mathbf{B}, \Sigma, \mathbf{Q} | \phi(f_h^{-1}(\mathbf{B}, \Sigma, \mathbf{Q})) > \mathbf{0})}{\int \pi(\mathbf{y}^T | \mathbf{B}, \Sigma) \pi(\mathbf{B}, \Sigma, \mathbf{Q} | \phi(f_h^{-1}(\mathbf{B}, \Sigma, \mathbf{Q})) > \mathbf{0}) d(\mathbf{B}, \Sigma, \mathbf{Q})}, \end{aligned}$$

where $\pi(\mathbf{y}^T | \mathbf{B}, \boldsymbol{\Sigma})$ is the likelihood function and $\pi(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q} | \phi(f_h^{-1}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})) > \mathbf{0})$ 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 $(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})$. 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 $(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})$ but the likelihood function.

The truncated likelihood function in Equation (15) can be written as

$$\pi(\mathbf{y}^T | \mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \Phi(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}) = \frac{[\Phi(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}] \pi(\mathbf{y}^T | \mathbf{B}, \boldsymbol{\Sigma})}{\int [\Phi(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}] \pi(\mathbf{y}^T | \mathbf{B}, \boldsymbol{\Sigma}) d\mathbf{y}^T}. \quad (16)$$

But note that

$$\begin{aligned} & \int [\Phi(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}] \pi(\mathbf{y}^T | \mathbf{B}, \boldsymbol{\Sigma}) d\mathbf{y}^T = \\ & = \int [\Phi(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}] \left(\prod_{t=1}^T \pi(\mathbf{y}_t | \mathbf{x}_t, \mathbf{B}, \boldsymbol{\Sigma}) \right) d(\mathbf{y}_1 \dots \mathbf{y}_T) \\ & = \int [\tilde{\Phi}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \boldsymbol{\varepsilon}^v) > \mathbf{0}] \left(\prod_{t=1}^T \frac{\pi(g_h^{-1}(\boldsymbol{\varepsilon}_t; \mathbf{x}_t, f_h^{-1}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})) | \mathbf{x}_t, \mathbf{B}, \boldsymbol{\Sigma})}{v_{g_h}(g_h^{-1}(\boldsymbol{\varepsilon}_t; \mathbf{x}_t, f_h^{-1}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})))} \right) d(\boldsymbol{\varepsilon}_1 \dots \boldsymbol{\varepsilon}_T), \end{aligned}$$

where $\tilde{\Phi}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \boldsymbol{\varepsilon}^v) = \phi(f_h^{-1}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}), \boldsymbol{\varepsilon}^v)$ and the term v_{g_h} is called the volume element of the function g_h evaluated at $g_h^{-1}(\boldsymbol{\varepsilon}_t; \mathbf{x}_t, f_h^{-1}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}))$. Our Equation (12) implies that

$$v_{g_h}(g_h^{-1}(\boldsymbol{\varepsilon}_t; \mathbf{x}_t, f_h^{-1}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}))) = |\boldsymbol{\Sigma}|^{-\frac{1}{2}} \text{ for } 1 \leq t \leq T.$$

Hence,

$$\begin{aligned} & \int [\tilde{\Phi}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \boldsymbol{\varepsilon}^v) > \mathbf{0}] \left(\prod_{t=1}^T \frac{\pi(g_h^{-1}(\boldsymbol{\varepsilon}_t; \mathbf{x}_t, f_h^{-1}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})) | \mathbf{x}_t, \mathbf{B}, \boldsymbol{\Sigma})}{v_{g_h}(g_h^{-1}(\boldsymbol{\varepsilon}_t; \mathbf{x}_t, f_h^{-1}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})))} \right) d(\boldsymbol{\varepsilon}_1 \dots \boldsymbol{\varepsilon}_T) \\ & = \int [\tilde{\Phi}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \boldsymbol{\varepsilon}^v) > \mathbf{0}] \left(\prod_{t=1}^T \pi(\boldsymbol{\varepsilon}_t) \right) d(\boldsymbol{\varepsilon}_1 \dots \boldsymbol{\varepsilon}_T) \\ & = \int [\tilde{\Phi}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \boldsymbol{\varepsilon}^v) > \mathbf{0}] \left(\prod_{s=1}^v \pi(\boldsymbol{\varepsilon}_{t_s}) \right) d(\boldsymbol{\varepsilon}_{t_1} \dots \boldsymbol{\varepsilon}_{t_v}). \end{aligned} \quad (17)$$

Equation (17) allows us to write the truncated likelihood in Equation (16) as

$$\pi(\mathbf{y}^T | \mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \Phi(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}) = \frac{[\Phi(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}] \pi(\mathbf{y}^T | \mathbf{B}, \boldsymbol{\Sigma})}{\omega(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})}, \quad (18)$$

where $\omega(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}) = \int [\tilde{\Phi}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \boldsymbol{\varepsilon}^v) > \mathbf{0}] (\prod_{s=1}^v \pi(\boldsymbol{\varepsilon}_{t_s})) d(\boldsymbol{\varepsilon}_{t_1} \dots \boldsymbol{\varepsilon}_{t_v})$. Equation (18) 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. Using Equation (18), the posterior of $(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})$ subject to the narrative sign restrictions is proportional to

$$\pi(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q} | \mathbf{y}^T, \Phi(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}) \propto \frac{[\Phi(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}] \pi(\mathbf{y}^T | \mathbf{B}, \boldsymbol{\Sigma})}{\omega(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})} \pi(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}). \quad (19)$$

One would normally choose priors of $(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})$ that are agnostic, i.e., uniform over $O(n)$.² When that is the case, Equation (19) can be written as

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

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, it can be shown that under agnostic priors the posterior distribution subject to the traditional sign restrictions is

$$\pi(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q} | \mathbf{y}^T, \phi(f_h^{-1}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})) > \mathbf{0}) \propto [\phi(f_h^{-1}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})) > \mathbf{0}] \pi(\mathbf{y}^T | \mathbf{B}, \boldsymbol{\Sigma}) \pi(\mathbf{B}, \boldsymbol{\Sigma}).$$

in which no re-weighting of the likelihood is needed.

If one uses both traditional and narrative sign restrictions, under agnostic priors, the posterior

²See Arias et al. (2016b) for a detailed treatment of agnostic priors. Suffice to say here that, given we only consider sign restrictions, agnostic priors of $(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})$ imply an agnostic prior of $\boldsymbol{\Theta}$ and that Algorithm 1, to follow, makes independent draws from the conditionally agnostic posterior of $\boldsymbol{\Theta}$ subject to the traditional sign restrictions.

distribution subject to the traditional and narrative sign restrictions is

$$\begin{aligned} & \pi(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q} | \mathbf{y}^T, \phi(f_h^{-1}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})) > \mathbf{0}, \Phi(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}) \\ & \propto [\phi(f_h^{-1}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})) > \mathbf{0}] \frac{[\Phi(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}, \mathbf{y}^v, \mathbf{x}^v) > \mathbf{0}] \pi(\mathbf{y}^T | \mathbf{B}, \boldsymbol{\Sigma})}{\omega(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})} \pi(\mathbf{B}, \boldsymbol{\Sigma}). \end{aligned} \quad (20)$$

We are now ready to specify our algorithm to draw from the posterior distribution defined in Equation (20).

4.2 The algorithm

The algorithm consists of three blocks and will allow us to combine traditional and narrative sign restrictions. The first block is identical to the algorithm in Rubio-Ramirez et al. (2010) and Arias et al. (2016b) since we use it to draw from the set of structural parameters satisfying the traditional sign restrictions.

Algorithm 1. *The following algorithm makes independent draws from the conditionally agnostic posterior of $(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})$ subject to the traditional sign restrictions.*

1. *Independently draw $(\mathbf{B}, \boldsymbol{\Sigma})$ from the posterior distribution of the reduced-form parameters.*
2. *Independently draw \mathbf{Q} from the uniform distribution over $O(n)$.*
3. *Check whether $\phi(f_h^{-1}(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})) > \mathbf{0}$.*
4. *If the restrictions are not satisfied, discard the draw of $(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})$. Otherwise, record the draw.*
5. *Return to Step 1 until the required number of draws has been obtained.*

There are various algorithms for making independent draws from the posterior of the reduced-form parameters, depending on the choice of prior over the reduced-form parameters. A common choice is to use conjugate priors. If the specified prior of the reduced-form parameters is conjugate,

then the posterior is normal-inverse-Wishart. This is a good choice because it is extremely easy and efficient to make independent draws from the normal-inverse-Wishart distribution.³

In order to draw from the posterior subject to the traditional and narrative sign restriction, we need to compute the weights $\omega(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$. The following algorithm can be used to approximate it.

Algorithm 2. *Given $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$, the following algorithm can be used to approximate $\omega(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$.*

1. *Independently draw $\boldsymbol{\varepsilon}^v$ from a Gaussian distribution with mean zero and covariance matrix $\mathbf{I}_{(n \times v)}$.*
2. *Check whether $\tilde{\Phi}(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}, \boldsymbol{\varepsilon}^v) > \mathbf{0}$.*
3. *Return to Step 1 until the required number of draws has been obtained.*

Algorithm 2 can be used to approximate the value of $\omega(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ by the share of draws of $\boldsymbol{\varepsilon}^v$ that satisfy $\tilde{\Phi}(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}, \boldsymbol{\varepsilon}^v) > \mathbf{0}$. It is important to take sufficient draws of Algorithm 2 to accurately approximate the importance weights. In practice, the larger v is, the more draws will be required.⁴

Finally, the algorithm to draw from the posterior distribution conditional on the traditional and narrative sign restrictions is

Algorithm 3. *The following algorithm makes independent draws from the conditionally agnostic posterior of $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ subject to the traditional and narrative sign restrictions.*

1. *Independently draw from the posterior of $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ subject to the traditional sign restrictions using Algorithm 1.*
2. *Check whether $\Phi(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}, \boldsymbol{\varepsilon}^v) > \mathbf{0}$.*

³For instance, Matlab implements both the normal and the inverse-Wishart distributions. Rubio-Ramirez et al. (2010) describe how to use the QR decomposition to independently draw the uniform distribution over the set of orthogonal matrices.

⁴In our empirical applications, 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.

3. *If a draw does not satisfy the narrative sign restrictions, set its importance weight to zero. Otherwise, use Algorithm 2 to approximate $\omega(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$, and set its importance weight to $\frac{1}{\omega(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})}$.*
4. *Return to Step 1 until the required number of draws that satisfy both the traditional and the narrative sign restrictions has been obtained.*
5. *Draw with replacement from the resulting set of $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ using the importance weights calculated in Step 3.*

Algorithm 3 makes clear that we need to importance weight the draws of $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ coming from Algorithm 1 by $\frac{1}{\omega(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})}$ in order to draw from the conditionally agnostic posterior of $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ subject to the traditional and narrative sign restrictions. If we do not importance sample the draws of $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ we will be making independent draws from a conditional posterior of $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ subject to the traditional and narrative sign restrictions but not from the conditionally agnostic posterior of $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ subject to the traditional and narrative sign restrictions implied by our agnostic prior.

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, consistent with the conclusions of Kilian and Murphy (2014).

After describing our data and baseline specification, we will report the list of historical events that we will use and the narrative sign restrictions they imply. We will first report the results associated with adding the whole list of events. Later, we will show that, in fact, adding a single narrative sign restriction by which the structural parameters must imply that an expansionary aggregate demand shock was not the main cause of the unexpected increase in the real price of oil observed in August 1990 is enough to obtain our results.

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.⁵ We extend their data set backward to January 1971 and forward to December 2015.

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 a negative oil supply shock leads to a decrease in oil production growth and economic activity, and an increase in the real price of oil; a positive aggregate demand shock leads to higher oil production growth, higher economic activity, and a higher real price of oil; and a positive oil-specific demand shock leads to higher oil production growth, lower economic activity, and a higher real price of oil. The restrictions on $\mathbf{L}_0(\Theta)$ are given in Table 1.

⁵The VAR is estimated on monthly data using 24 lags and a constant. 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. This specification has proved successful in out-of-sample forecasting and is the most widely used in the literature. For some alternatives, see Juvenal and Petrella (2015) or Baumeister and Hamilton (2015).

Kilian and Murphy (2012) 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, i.e., $(\mathbf{L}_0(\Theta))_{1,2} / (\mathbf{L}_0(\Theta))_{3,2}$ and $(\mathbf{L}_0(\Theta))_{1,3} / (\mathbf{L}_0(\Theta))_{3,3}$. 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 price elasticity of supply restriction as the baseline specification.

Table 1: SIGN RESTRICTIONS ON IMPACT RESPONSES

	Oil Supply Shock	Aggregate Demand Shock	Oil-specific Demand Shock
Oil Production Growth	−	+	+
Economic Activity Index	−	+	−
Real Oil Price	+	+	+

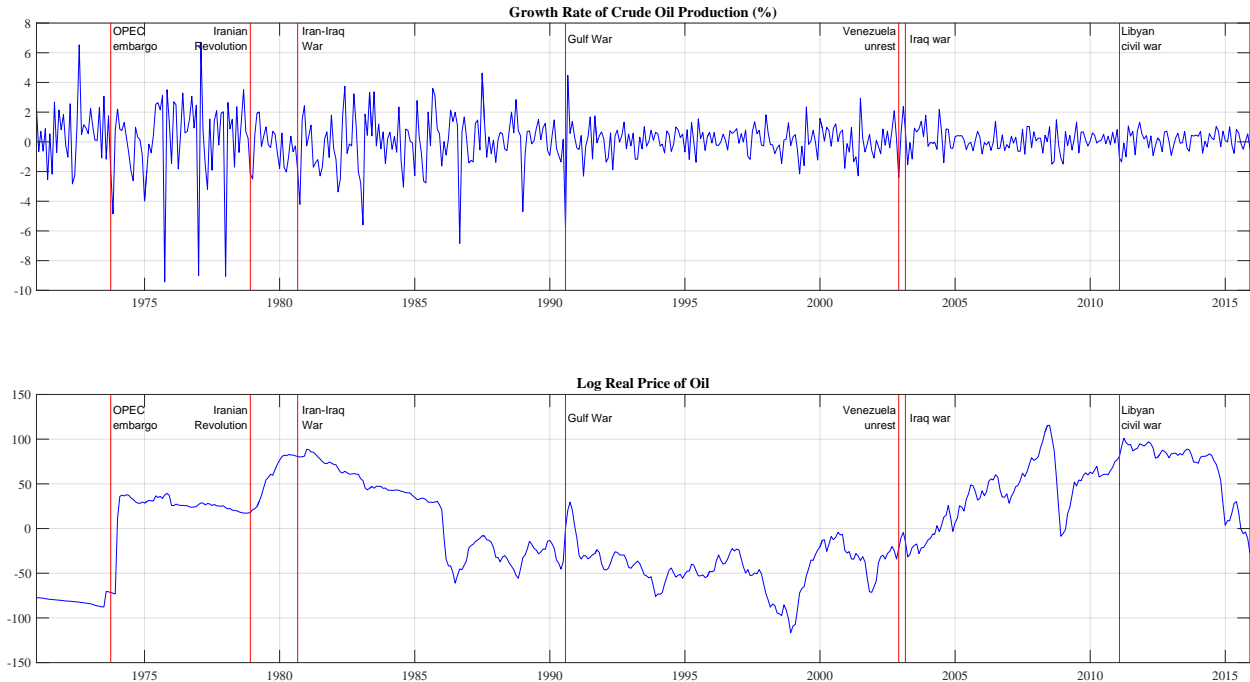
5.2 The narrative information

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 red lines: the Yom Kippur War and subsequent OPEC embargo (October 1973), the start of the Iranian Revolution (December 1978-January 1979), the outbreak of the Iran-Iraq War (September-October 1980), the Iraqi invasion of Kuwait, which marked 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).⁶ It is obvious that these historical events had a major impact both

⁶The 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. The Libyan Civil War erupted in February 2011 in the context of wider protests in

on the production growth and the real price of oil. The fact that each of these historical events is exogenous and oil production growth was negatively affected makes them clear candidates for negative oil supply shocks.

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 OPEC 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).

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 “no OPEC oil facilities were attacked during the October war, and there is no evidence of OPEC production shortfalls caused by military action” (Kilian, 2008, p. 218) citing the role of global demand and inflationary pressures in commodity markets as a potential driver of the increase in the price of oil. 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.

oil during that episode. Since there is no agreement on this particular event, we will therefore exclude the 1973 episode.⁷ Thus, we will therefore impose the following narrative sign restriction:

Narrative Sign Restriction 1. *The oil supply shock must take negative values in December 1978, January 1979 (outbreak of the Iranian Revolution), in September and October 1980 (outbreak of the Iran-Iraq War), August 1990 (outbreak of the Persian Gulf War), December 2002 (Venezuela oil strike), March 2003 (outbreak of the Iraq War) and February 2011 (outbreak of the Libyan Civil War).*

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

⁷Moreover, 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.

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 of the data into structural shocks.

5.3 Results

Figure 2 displays IRFs of the three variables to the three structural shocks, with and without the narrative information. The gray shaded area represents the 68% (point-wise) highest posterior density (HPD) credible sets for the IRFs and the solid blue lines are the point-wise median IRFs using the baseline identification.⁸ The pink shaded areas and red solid lines display the equivalent quantities

⁸It 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.

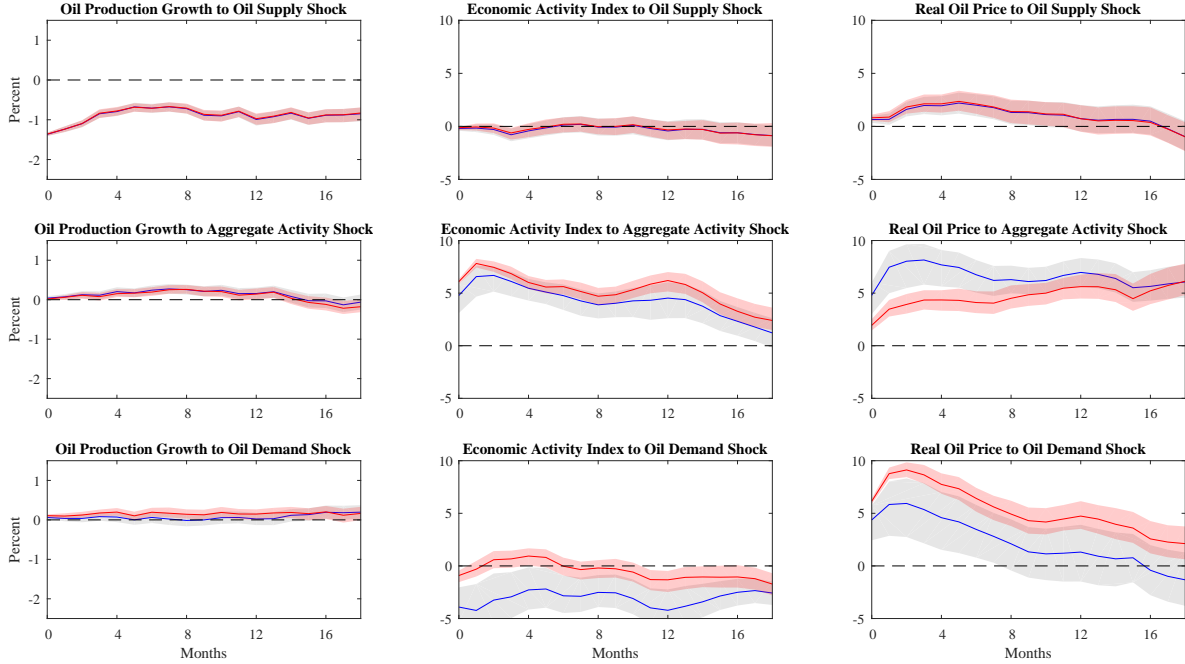
when Narrative Sign Restrictions 1-3 are also used.⁹ The narrative sign restrictions dramatically narrow down the uncertainty around many of the IRFs relative to the baseline identification and modifies the shape of some of the IRFs in economically meaningful ways. Oil-specific demand shocks are shown 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 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 a positive oil-specific demand shock, somewhat counter-intuitively, caused a temporary boom in economic activity.¹⁰

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

⁹Narrative 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 of Algorithm 2 above.

¹⁰The results using the Cholesky decomposition that we refer to can be seen in Figure 3, p. 1061, in Kilian (2009b).

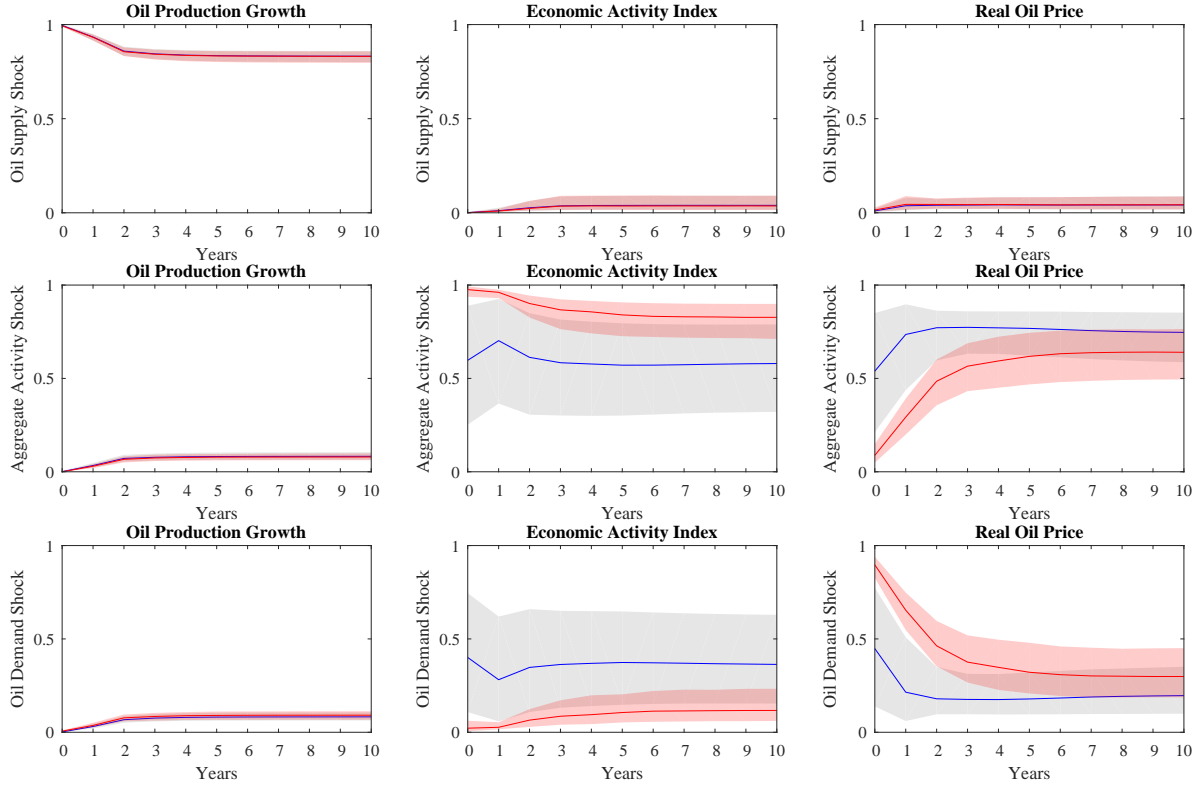
Figure 2: IRFs WITH AND WITHOUT NARRATIVE SIGN RESTRICTIONS



Note: The gray shaded area represents the 68% (point-wise) HPD credible sets for the IRFs, and the solid blue lines are the median IRFs using the baseline identification restrictions. The pink shaded areas and red 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.

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 gray and the pink shaded areas we see that adding the narrative sign restrictions (pink 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 helps disentangle the effects of aggregate demand and oil-specific

Figure 3: FEVD WITH AND WITHOUT NARRATIVE SIGN RESTRICTIONS

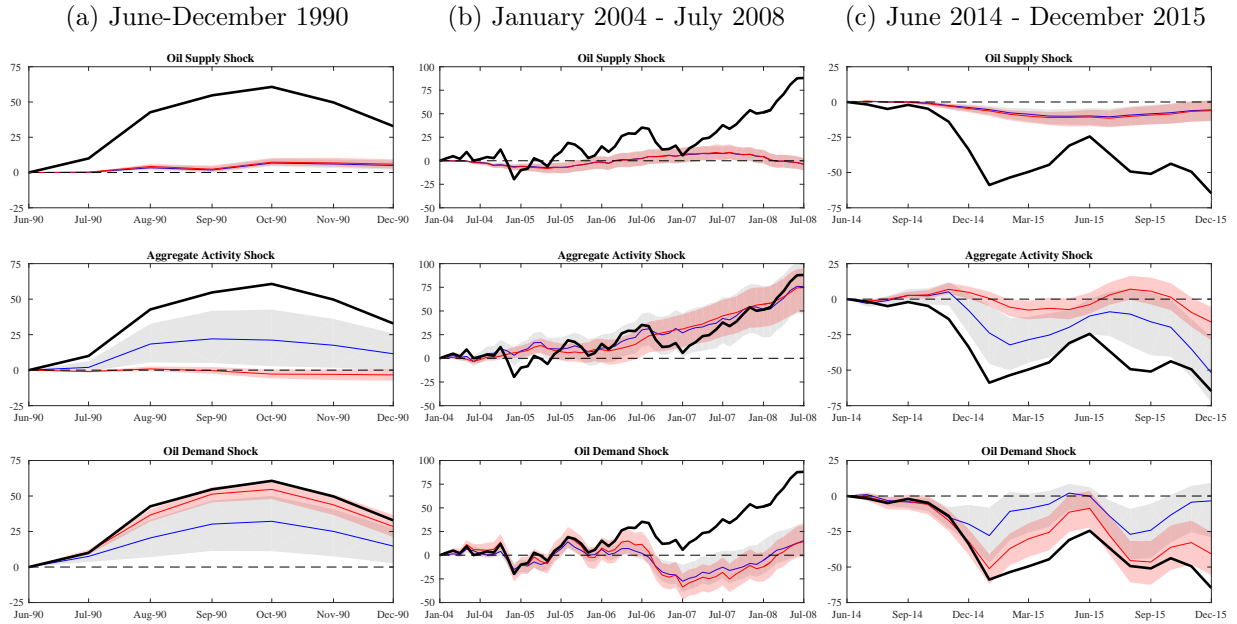


Note: Each panel presents the estimated contribution of each structural shock to the mean squared forecast error at horizons of 1-10 years for the three variables, expressed as a percentage of the total MSE. The gray shaded area represents the 68% (point-wise) HPD credible sets for the FEVD, and the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

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 (gray shaded area) is consistent with many structural parameters that imply that aggregate demand shocks were important contributors to the unexpected increase in

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 black line. The solid blue lines are the median for the baseline identification restrictions, while the gray shaded area represents the 68% (point-wise) HPD credible sets. The red solid lines and the pink shaded areas display the equivalent quantities when Narrative Sign Restrictions 1-3 are also satisfied.

log real oil prices observed between July and November 1990. Including Narrative Sign Restrictions 1-3 (pink shaded area) reinforces the view of Kilian (2009b) and 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 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 clearly point toward oil-specific demand shocks as the source of the collapse.¹¹

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 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. 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. In other words, one only needs to agree that expansionary

¹¹Baumeister and Kilian (2016) and Kilian (2016) have studied the same period using an alternative methodology.

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%

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. To see this more clearly, we can consider Alternative 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.*

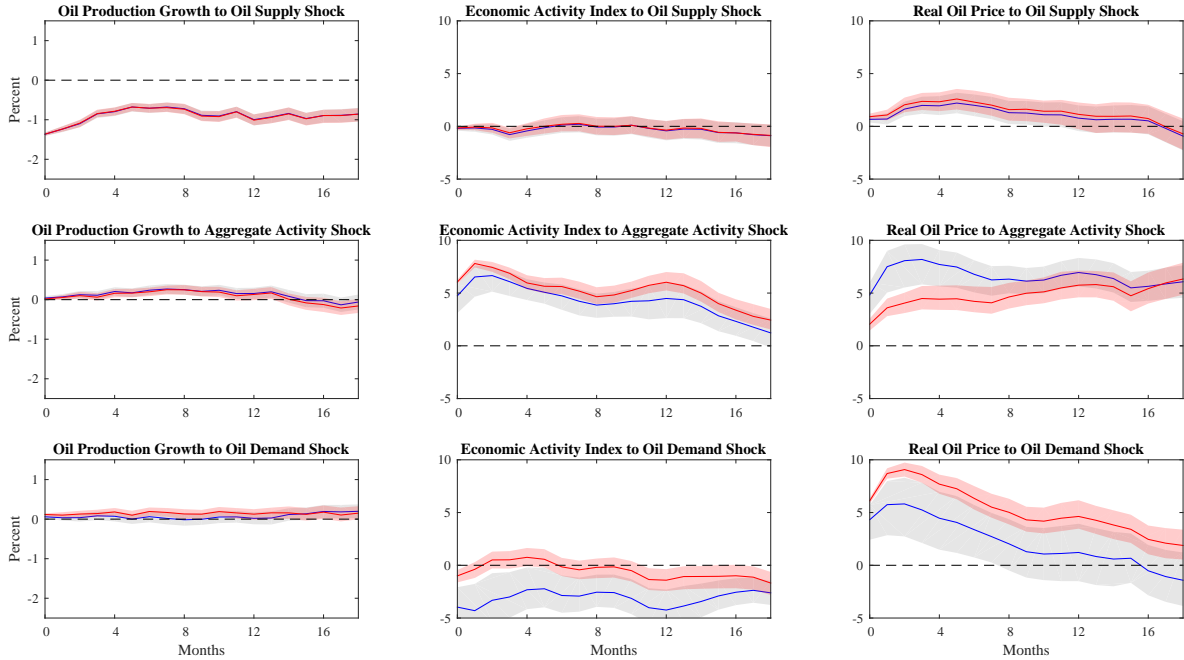
Figure 5 plots the same IRFs reported in Figure 2, but the pink shaded areas and red solid lines now add the Alternative Narrative Sign Restriction 3 to the baseline identification instead of adding the Narrative Sign Restrictions 1-3.¹² As the reader can see, Figures 2 and 5 are almost identical.¹³ Hence using either set of narrative sign restrictions has comparable effects on the IRFs

¹²Alternatively, 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.

¹³Alternative Narrative Sign Restriction 3 affects in total one time period. Ten thousand draws that satisfy the

Figure 5: IRFs WITH AND WITHOUT NARRATIVE SIGN RESTRICTIONS

(ALTERNATIVE NARRATIVE SIGN RESTRICTION 3)



Note: The gray shaded area represents the 68% (point-wise) HPD credible sets for the IRFs, and the solid blue lines are the median IRFs using the baseline identification restrictions. The pink shaded areas and red 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.

and on other results such as the FEVD and historical decompositions presented above.¹⁴ 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

baseline restrictions are generated. Out of these, 749 satisfy the narrative sign restrictions. We approximate their weight in the importance step by using one million draws of Algorithm 2 above.

¹⁴The 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.

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.

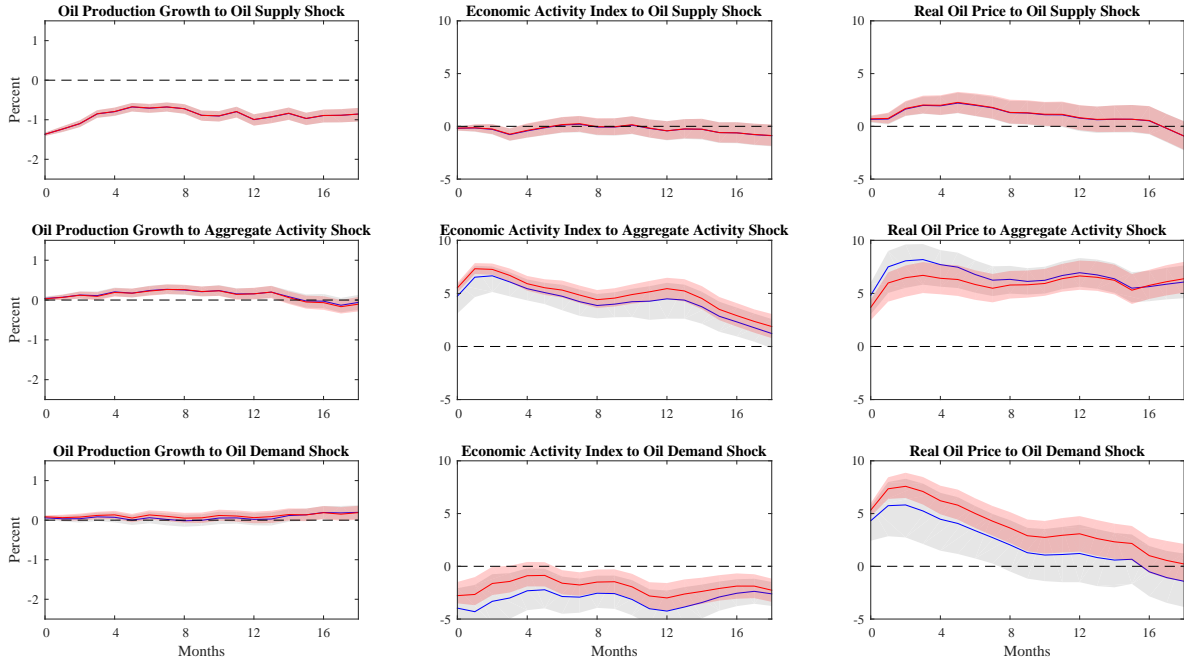
Given that the restriction relating to August 1990 appears to be key to our results, it warrants some additional discussion. In particular, we will analyze the robustness of the results to using the Type B variant of Alternative Narrative Sign Restriction 3, instead of the Type A variant we have been using so far. Recall from Section 3.4 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 3.4.3). Figure 6 plots the same IRFs reported in Figure 5 when adding the Alternative Narrative Sign Restriction 3 to the baseline identification, but in its milder Type B variant.¹⁵ As the reader can see, the main conclusions are maintained. 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.

5.5 Final remarks on demand and supply shocks in the oil market

To sum up, we have shown that, while the baseline identification proposed by Kilian and Murphy (2012) does a very good job of distinguishing the effects of supply and demand shocks, the narrative sign restrictions are successful in sharpening the inference about the structural shocks that drive the oil market. Using Kilian (2008) and Hamilton (2009) as sources, we obtain a list of post-1973

¹⁵Alternative Narrative Sign Restriction 3 (Type B) affects in total two time periods. Ten thousand draws that satisfy the baseline restrictions are generated. Out of these, 4500 satisfy the narrative sign restriction. We approximate their weight in the importance step by using one thousand draws of Algorithm 2 above.

Figure 6: IRFs WITH AND WITHOUT SIGN NARRATIVE RESTRICTIONS
 (ALTERNATIVE NARRATIVE SIGN RESTRICTION 3 – TYPE B)



Note: The gray shaded area represents the 68% (point-wise) HPD credible sets for the IRFs and the solid blue lines are the median IRFs using the baseline identification restrictions. The pink shaded areas and red 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.

historical events that generate a number of uncontroversial narrative sign restrictions that allow us to distinguish between aggregate demand and oil-specific demand shocks. In fact, it turns out that a single narrative sign restriction that ensures that the structural parameters are in line with the established narrative about the outbreak of the Persian Gulf War, whether in its stronger or milder variant, is enough to separate the effects of these two shocks. The fact that a single sign narrative restriction is enough is very important given that we started this section with a query for few uncontroversial sign restrictions that may help us reduce the set of structural parameters consistent with the baseline identification.

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, 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) among others, 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.

As we will see, once we add to Uhlig’s (2005) traditional sign restrictions our narrative sign

Table 3: SIGN RESTRICTIONS ON RESPONSES AT HORIZONS 0 TO 5

Monetary Policy Shock	
Real GDP	
GDP Deflator	–
Commodity Price Index	–
Total Reserves	–
Nonborrowed Reserves	–
Federal Funds Rate	+

restrictions such that only structural parameters that imply an important negative monetary policy shock occurred in October of 1979 are retained, the “consensus” revives. Given that the challenge is to come up with few additional uncontentious sign restrictions that help shrink the set of admissible structural parameters, the fact that the “consensus” is recovered by just considering narrative information about a single event is a great achievement.

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, to maximize comparability with previous work we chose the exact specification, reduced-form prior and data definitions used in the aforementioned papers. Our sample period is January 1965 to November 2007.¹⁶ Our baseline identification is identical to Uhlig’s (2005). Specifically, he postulates that a contractionary monetary policy shock has the effects described in Table 3 for the first six months.

¹⁶The VAR is estimated on monthly data using 12 lags and no constant or deterministic trends. We refer to the aforementioned papers for details on the sources and 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.

6.2 The narrative information

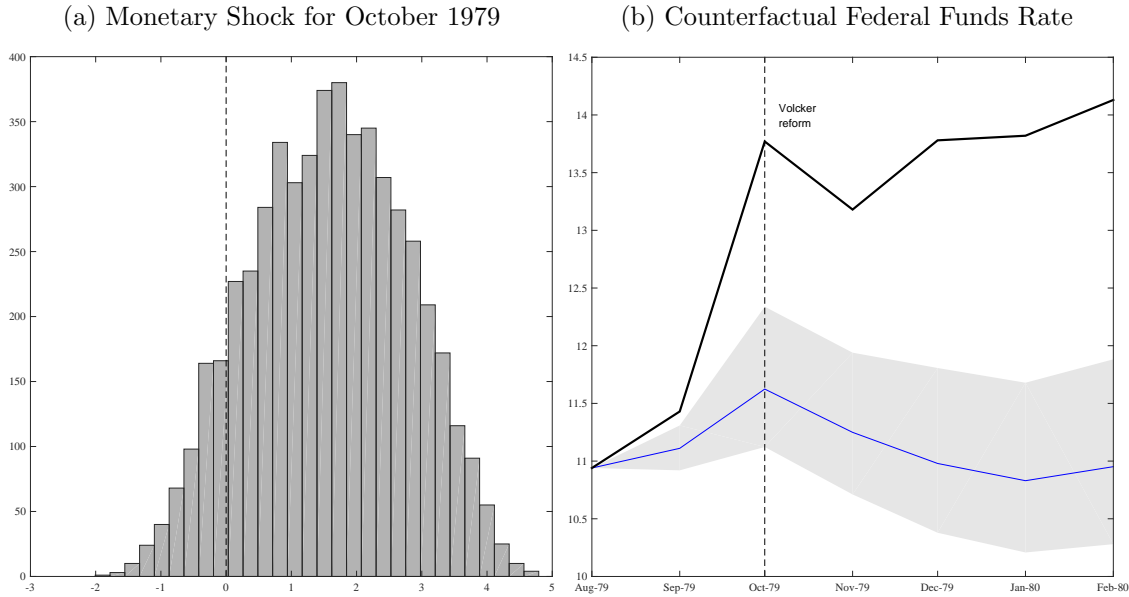
We start by examining the implications of the baseline specification for the period around October 1979. Panel (a) of Figure 7 displays the posterior distribution of the monetary policy shock during that month. While most of the distribution has positive support (i.e., a contractionary monetary policy 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 (blue line with gray 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 October. So the baseline specification effectively implies that the tightening 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 will 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 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.*

Figure 7: RESULTS AROUND OCTOBER 1979 WITH BASELINE IDENTIFICATION



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

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 of the fed funds rate into structural shocks.

Note that the specified Type B restriction postulates that the absolute value of the contribution of the monetary policy shock is “larger than the sum of the absolute value of the contribution of all other structural shocks” to the unexpected movement in the federal funds rate in October 1979,

whereas 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:

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.*

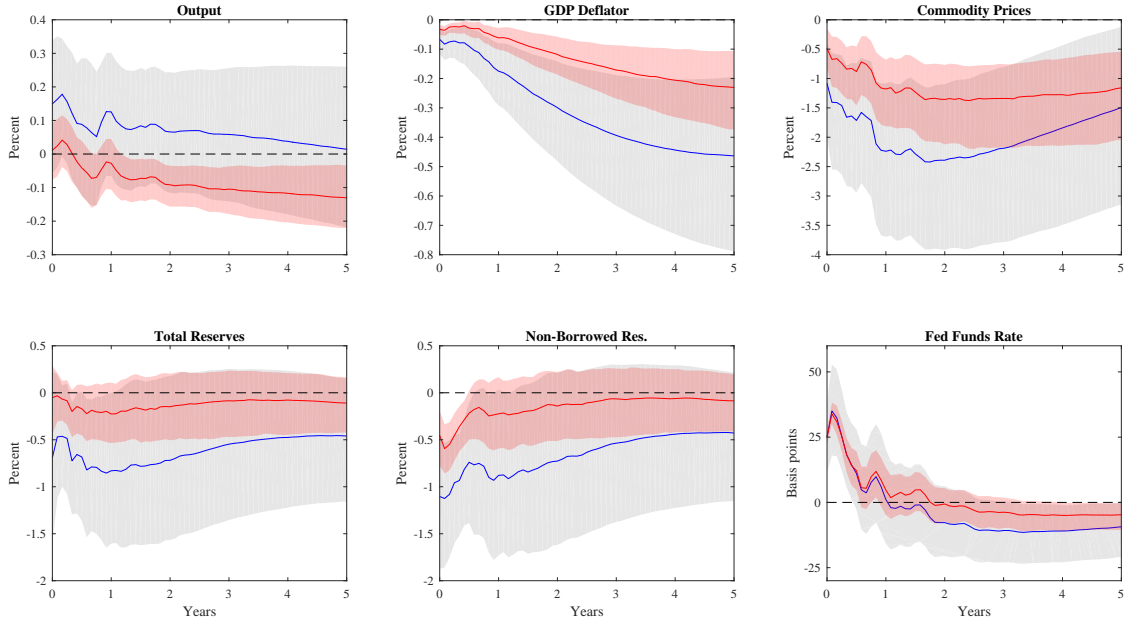
6.3 Results

Figure 8 compares the IRFs to a monetary policy shock, with and without narrative sign restrictions. The gray shaded area represents the 68% (point-wise) HPD credible sets for the IRFs and the solid blue lines are the median IRFs using the baseline identification. These results replicate the IRFs depicted in Figure 6 of Uhlig (2005). The pink shaded areas and red solid lines display the equivalent quantities when Narrative Sign Restrictions 4 and 5 are also used.¹⁷ 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

¹⁷Narrative 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 of Algorithm 2 above.

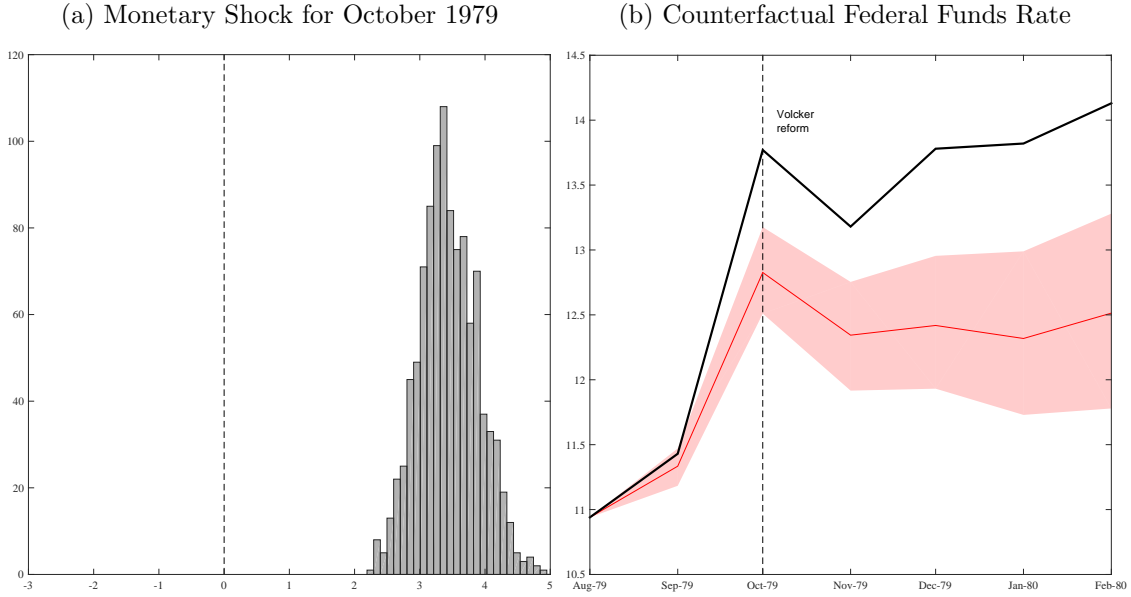
Figure 8: IRFs WITH AND WITHOUT NARRATIVE SIGN RESTRICTIONS



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

October 1979? Figure 9 plots the same results displayed in Figure 7, but this time with the narrative sign restrictions in place. Panel (a) of Figure 9 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 (red line with pink 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. The monetary policy shock was the overwhelming contributor to the unexpected increase in the federal funds rate. The figure tells us 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

Figure 9: RESULTS AROUND OCTOBER 1979 WITH NARRATIVE SIGN RESTRICTIONS



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

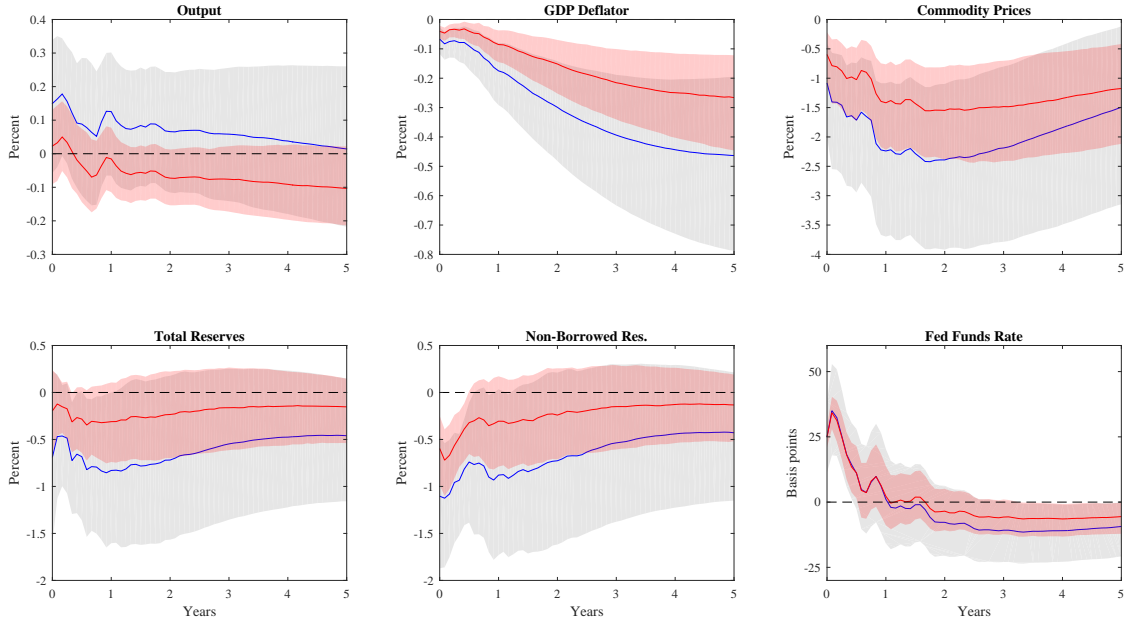
Narrative Sign Restrictions 4 and 5; only the sign of the shock and the sign of the contribution of the monetary policy shock relative to other structural shocks are.

As mentioned above, Narrative Sign Restriction 5 is the strongest version of the restriction. Figure 10 displays the main results when the milder Alternative Narrative Sign Restriction 5 is used instead.¹⁸ With the weaker restrictions the HPD credible sets are wider, but the basic message survives: output drops after a monetary policy shock with very high posterior probability. 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

¹⁸Narrative 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 of Algorithm 2 above.

Figure 10: IRFs WITH AND WITHOUT NARRATIVE SIGN RESTRICTIONS

(NARRATIVE SIGN RESTRICTION 4 AND ALTERNATIVE NARRATIVE SIGN RESTRICTION 5)

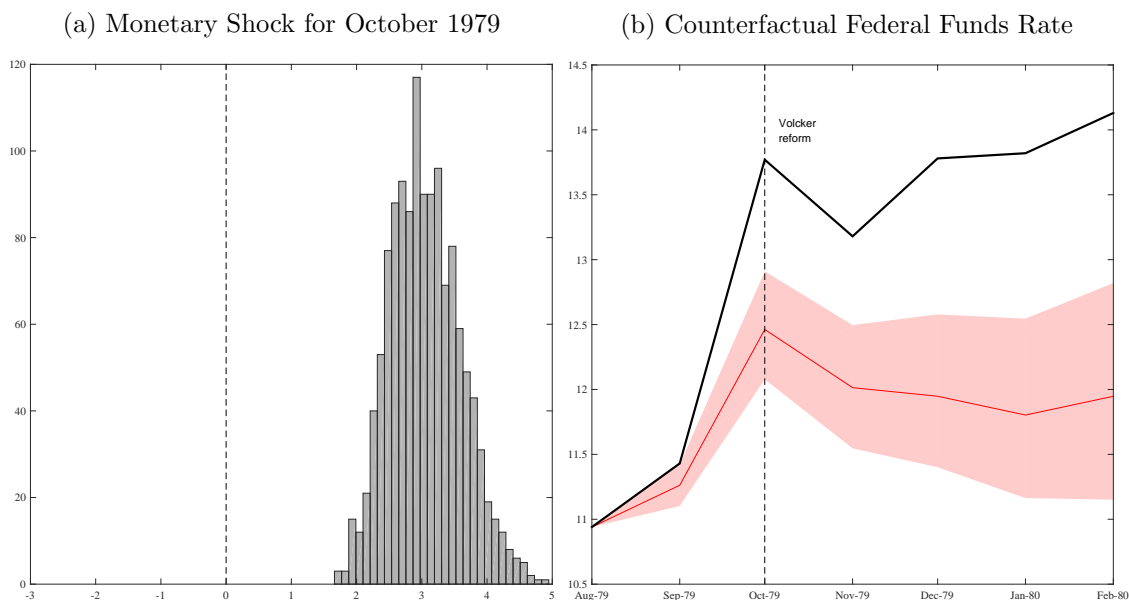


Note: The gray shaded area represents the 68% (point-wise) HPD credible sets for the IRFs, and the solid blue lines are the median IRFs using the baseline identification restrictions. The pink shaded areas and red solid lines display the equivalent quantities for the models that additionally satisfy 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.

conclude that monetary policy shocks reduce output. Our results echo those of Inoue and Kilian (2013) and Arias et al. (2016a), which question the robustness of Uhlig’s (2005) introduction of additional restrictions.

Alternative Narrative Sign Restriction 5 does not meaningfully change the implications for the period around October 1979 relative to Narrative Sign Restriction 5. Figure 11 replicates the panels displayed in Figure 9, 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

Figure 11: RESULTS AROUND OCTOBER 1979 WITH NARRATIVE SIGN RESTRICTIONS
 (NARRATIVE SIGN RESTRICTION 4 AND ALTERNATIVE NARRATIVE SIGN RESTRICTION 5)



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

October of 1979.

6.4 Including additional events: a new chronology

The results above have highlighted that using narrative information for a single event – October 1979 – is enough to alter the effect of monetary policy shocks on output reported in Uhlig (2005). That event is in our view the clearest and most uncontroversial example of a monetary policy shock, but as mentioned above, there is a long literature that uses historical sources to isolate monetary policy shocks. In Appendix A, we survey this literature and drawing, in particular, on Romer and Romer (1989), Romer and Romer (2004), and Gürkaynak et al. (2005), we identify eight events for

which there appears to be reasonable agreement that a monetary policy shock occurred. Of these, four – April 1974, October 1979, December 1988 and February 1994 – were contractionary shocks and four – December 1990, October 1998, April 2001, and November 2002 – were expansionary shocks. The inclusion of narrative information for all these events leads to results very similar to those reported above using only October 1979 but, as expected, the HPD credible sets narrow further.¹⁹ In particular, Figure A.2 shows the IRFs when Narrative Sign Restrictions 4 and 5 are used together with the following additional ones:

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

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.*

Narrative Sign Restriction 6 is equivalent to Narrative Sign Restriction 4, while Narrative Sign Restriction 7 is equivalent to Alternative Narrative Sign Restriction 5 for the additional dates. This implies that we impose only the weaker Type A restriction for all events other than the Volcker Reform. As mentioned above, this is because the Volcker Reform is the clearest episode.

Finally, we also would like to mention that while adding the narrative sign restrictions relating to the Volcker Reform are sufficient to obtain our results, they are not necessary. It is possible to

¹⁹Narrative Sign Restrictions 4-7 affect in total eight structural shocks. Five hundred thousand draws that satisfy the baseline restrictions are generated. Out of these, 553 satisfy the narrative sign restrictions. We approximate their weight in the importance step by using one million draws of Algorithm 2 above.

obtain results similar to the ones reported in Figures 8 and 10 by just imposing equivalent narrative restrictions only for the 1998, 1994 or 2001 dates on their own.²⁰

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.

²⁰These additional figures are available upon request.

- 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 J. D. HAMILTON (2015): “Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks,” Mimeo.
- BAUMEISTER, C. AND L. KILIAN (2016): “Understanding the Decline in the Price of Oil since June 2014,” *Journal of the Association of Environmental and Resource Economists*, 3, 131 – 158.
- 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.
- CALDARA, D. AND E. HERBST (2016): “Monetary Policy, Real Activity, and Credit Spreads: Evidence from Bayesian Proxy SVARs,” FEDs Working Paper Series 2016-049, Board of Governors of the Federal Reserve System (U.S.).
- 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.
- DRAUTZBURG, T. (2016): “A Narrative Approach to a Fiscal DSGE Model,” Working paper 16-11, Federal Reserve Bank of Philadelphia.

- 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.
- HOOVER, K. D. AND S. J. PEREZ (1994): “Post hoc ergo propter hoc once more: an evaluation of ‘does monetary policy matter?’ in the spirit of James Tobin,” *Journal of Monetary Economics*, 34, 47–74.
- INOUE, A. AND L. KILIAN (2013): “Inference on impulse response functions in structural VAR models,” *Journal of Econometrics*, 177, 1–13.
- JUVENAL, L. AND I. PETRELLA (2015): “Speculation in the oil market,” *Journal of Applied Econometrics*, 30, 621–649.
- 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,” in *Causes and Consequences of the Oil Shock of 2007-08*, ed. by J. Hamilton, vol. 40, 215–283.
- (2009b): “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market,” *American Economic Review*, 99, 1053–69.
- (2016): “The Impact of the Fracking Boom on Arab Oil Producers,” CEPR Discussion Papers 11107, C.E.P.R. Discussion Papers.
- 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*.

- 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.
- RAMEY, V. A. AND M. D. SHAPIRO (1998): “Costly capital reallocation and the effects of government spending,” *Carnegie-Rochester Conference Series on Public Policy*, 48, 145–194.
- 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.
- 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 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). The Romers’ 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.

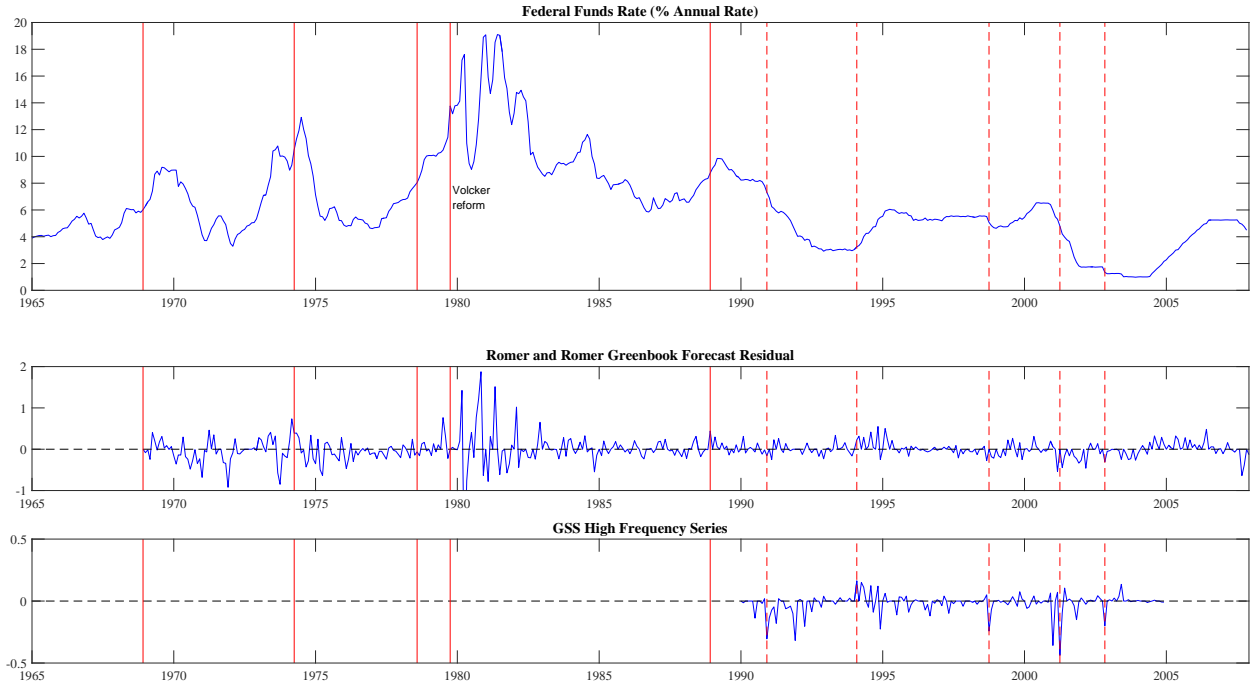
It is important to notice that previous approaches have proceeded by first constructing a time series of monetary policy shocks and then using that series either directly (e.g., Romer and Romer (1994)) or as an instrument (e.g., Gertler and Karadi (2015)) to estimate the effects on variables of interest. But owing to differences in methods and sources, the existing narrative series are sometimes inconclusive and other times contradictory, and one may not necessarily agree with the sign and the magnitude of every single observation in those series. Instead, we will draw on a variety of sources and previous narrative series to isolate a few events that arguably constitute clear monetary policy shocks, and impose that information directly as sign restrictions in the SVAR. Thus, our method is closer to the spirit of Friedman and Schwartz (1963) and Romer and Romer (1989), but inherits the advantages of sign-restricted SVARs.

The solid vertical lines in Figure A.1 represent the original 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. As mentioned above, during the subsamples in which the series overlap, they disagree a great deal. We combine the three approaches to select the dates for which the evidence of an exogenous monetary policy shock is most compelling.

For the first half of the sample, on which the Romers' original analysis was conducted, we revisit their original dates in light of the Greenbook series:

- *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

Figure A.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), whereas the dashed vertical lines represent the additional episodes identified in the chronology below.

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 OPEC 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.¹
- *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

¹Note 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.

four months. Therefore, the evidence appears to favor the occurrence of a monetary policy shock during this period.

It is worth it at this point to make two observations. The first is that the RR-04 series, like the actual fed funds rate, displays very large movements during the reserves-targeting period of October 1979 to late 1982, but these appear to be a side effect of the abandonment of the funds rate target, rather than shocks associated with any particular identifiable event. The second is that Hoover and Perez (1994) criticized the RR-89 chronology, pointing to the possibility that the oil price shocks of the 1970s were in fact causing the Romers' monetary policy shifts. Since our VAR specification includes lagged values both of inflation and commodity prices, the narrative restrictions we will impose will refer to the *unexpected* component of the federal funds rate, which is by construction unforecastable from past price developments. Therefore, this concern is alleviated when combining the VAR and narrative approaches.

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), 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 identify as candidate events December 1990, February 1994, October 1998, January 2001 and November 2002. 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.

- *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.²

- *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 came as a surprise to financial market participants, leading to a large adjustment in longer-term interest rates.³ 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 February 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

²The main justification for the surprisingly dovish stance appears to be an unwillingness to sacrifice output 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.”

³See “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.

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. We will therefore keep this event in the chronology and assess its importance for the results.

- *October 1998.* In late September of 1998, the FOMC responded to the deterioration in 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.”⁴ 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.”⁵ 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.
- *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

⁴See *Statement*, Federal Open Market Committee, September 29, 1998.

⁵See *Transcript*, Federal Open Market Committee, October 15, 1998.

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. 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, its 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. On the other hand, the Business Cycle Dating Committee of the NBER later declared that a recession had started in March 2001 so it could be the case that in moving in April, the FOMC was foreseeing further economic weakness. We will initially keep this event in the chronology but assess its importance for the results.

- *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.⁶ Geopolitical risks – preparations

⁶One 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 [...]

for the 2003 Iraq war were already under way – were also a concern.⁷ 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 points- it is possible that 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.”⁸

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 identified 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 following narrative restrictions:

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

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.*

we were looking into a deep deflationary hole.” See *Transcript*, Federal Open Market Committee, November 6, 2002.

⁷See *Statement*, Federal Open Market Committee, November 6, 2002.

⁸See *Transcript*, Federal Open Market Committee, November 6, 2002.

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. In addition, given that for the Volcker episode the evidence is stronger, we will also consider Narrative Sign Restriction 5.

A.1 Results

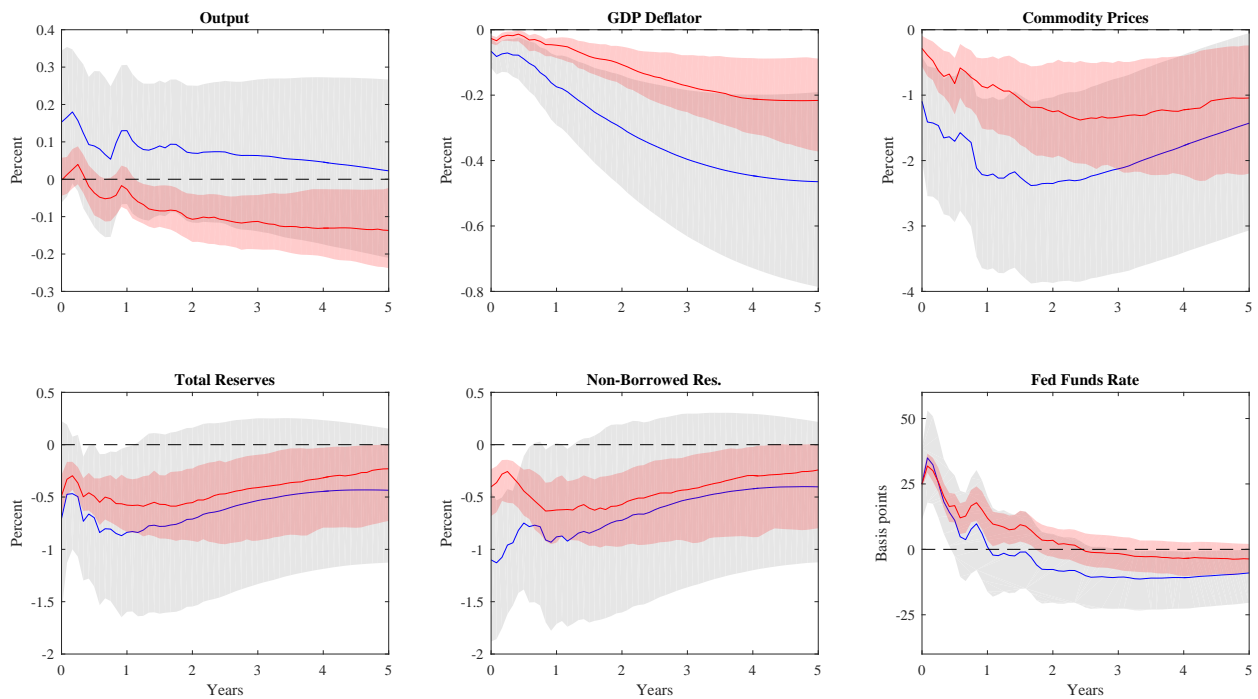
Figure A.2 presents the IRFs to a monetary policy shock, with and without narrative information. The gray shaded area represents the 68% (point-wise) HPD credible sets for the IRFs and the solid blue lines are the median IRFs using the baseline identification. These results replicate the IRFs depicted in Figure 6 of Uhlig (2005). The pink shaded areas and red solid lines display the equivalent quantities when Narrative Sign Restrictions 5-7 are also used. The results are very similar to those using only the Volcker episode, reported in Figure 8 in the main text.

Table A.1 looks at the probability that the baseline model violates the narrative restrictions, individually and jointly. Looking at the last column of the table, it can be seen that the baseline model disagrees with the narrative information for most of the episodes with a very high probability.

Table A.1: PROBABILITY OF VIOLATING THE NARRATIVE RESTRICTIONS

	Restr. 5	Restr. 6	Restr. 7	Any Restr.
April 1974	—	2.7%	50%	50%
October 1979	90%	11%	78%	90%
December 1988	—	71%	85%	96%
December 1990	—	2.8%	52%	52%
February 1994	—	77%	84%	97%
October 1998	—	48%	86%	86%
April 2001	—	55%	90%	91%
November 2002	—	44%	86%	86%
Any Episodes	—	96%	99.8%	99.9%

Figure A.2: IRFs WITH AND WITHOUT NARRATIVE SIGN RESTRICTIONS
 (NARRATIVE SIGN RESTRICTIONS 5, 6, AND 7)



Note: The gray shaded area represents the 68% (point-wise) HPD credible sets for the IRFs, and the solid blue lines are the median IRFs using the baseline identification restrictions. The pink shaded areas and red solid lines display the equivalent quantities for the models that additionally satisfy Narrative Sign Restrictions 5-7.