

NBER WORKING PAPER SERIES

DRAWING CONCLUSIONS FROM STRUCTURAL VECTOR AUTOREGRESSIONS
IDENTIFIED ON THE BASIS OF SIGN RESTRICTIONS

Christiane Baumeister
James D. Hamilton

Working Paper 26606
<http://www.nber.org/papers/w26606>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2020

The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Christiane Baumeister and James D. Hamilton. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Drawing Conclusions from Structural Vector Autoregressions Identified on the Basis of Sign Restrictions

Christiane Baumeister and James D. Hamilton

NBER Working Paper No. 26606

January 2020

JEL No. C30,E5,F2

ABSTRACT

This paper discusses the problems associated with using information about the signs of certain magnitudes as a basis for drawing structural conclusions in vector autoregressions. We also review available tools to solve these problems. For illustration we use Dahlhaus and Vasishtha's (2019) study of the effects of a U.S. monetary contraction on capital flows to emerging markets. We explain why sign restrictions alone are not enough to allow us to answer the question and suggest alternative approaches that could be used.

Christiane Baumeister
Department of Economics
University of Notre Dame
3028 Jenkins Nanovic Hall
Notre Dame, IN 46556
and NBER
cbaumeis@nd.edu

James D. Hamilton
Department of Economics, 0508
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0508
and NBER
jhamilton@ucsd.edu

Drawing structural conclusions on the basis of dynamic correlations requires identifying information with which to interpret the correlations observed in the data. Because many researchers have doubts about the validity of this identifying information, it has become a very common practice to use information about the signs of certain magnitudes as a basis for identification in structural vector autoregressions. Prominent examples include Aastveit, Bjørnland and Thorsrud (2015), Aastveit et al. (2017), Abbate et al. (2016), Abbate, Eickmeier and Prieto (2016), Abdallah and Lastrapes (2013), Amir-Ahmadi, Matthes and Wang (2016, 2017), Antolín-Díaz and Rubio-Ramírez (2018), Anzuini, Lombardi and Pagano (2013), Arias, Caldara and Rubio-Ramírez (2015), Arias, Rubio-Ramírez and Waggoner (2018), Baumeister and Peersman (2013a, 2013b), Benati (2015), Benati and Lubik (2014), Belongia and Ireland (2016), Berg (2012), Bian and Gete (2015), Bjørnland and Halvorsen (2014), Boeckx, Dossche and Peersman (2017), Breitenlechner, Scharler and Sindermann (2016), Buch, Eickmeier and Prieto (2014), Canova and De Nicoló (2002), Canova and Paustian (2011), Castelnovo and Surico (2010), Chadha, Corrado and Sun (2010), Charnavoki and Dolado (2014), Conti, Neri and Nobili (2017), Corsetti, Dedola and Leduc (2014), Dahlhaus and Vasishtha (2019), Darracq Paries and De Santis (2015), Dedola and Neri (2007), Eickmeier and Hofmann (2013), Eickmeier and Ng (2015), Ellis, Mumtaz and Zabczyk (2014), Enders, Müller and Scholl (2011), Faccini, Mumtaz and Surico (2016), Fadejeva, Feldkircher and Reininger (2017), Farrant and Peersman (2006), Feldkircher and Huber (2016), Fornari and Stracca (2012), Feroni, Furlanetto and Lepetit (2018), Fratzscher, Juvenal and Sarno (2010), Fratzscher and Straub (2013), Furlanetto, Ravazzolo and Sarferaz (2017), Furlanetto

and Robstad (2017), Fujita (2011), Gambacorta, Hofmann and Peersman (2014), Gambetti and Musso (2017), Georgiadis (2015), Glocker and Towbin (2015), Güntner (2014), Gupta and Modise (2013), Huber and Punzi (2017), Hristov, Hülsewig and Wollmershäuser (2012), Hofmann, Peersman and Straub (2012), Jääskelä and Jennings (2011), Jarociński (2010), Juvenal (2011), Juvenal and Petrella (2014), Kapetanios et al. (2012), Kilian and Murphy (2012, 2014), Kilian and Zhou (2019), Kim, Moon and Velasco (2017), Lippi and Nobili (2012), Liu et al. (2016), Liu et al. (2019), Liu, Mumtaz and Theophilopoulou (2014), Luciani (2015), Meeks (2012), Michaelis and Watzka (2017), Mönch (2012), Mountford (2005), Mountford and Uhlig (2009), Mumtaz and Sunder-Plassmann (2013), Mumtaz and Zanetti (2012), Pappa (2009), Peersman (2005), Peersman and Straub (2009), Riggi and Venditti (2015), Rubio-Ramírez, Waggoner and Zha (2010), Sá, Towbin and Wieladek (2014), Sá and Wieladek (2015), Schenkelberg and Watzka (2013), Scholl and Uhlig (2008), Uhlig (2005), Van Robays (2016), Vargas-Silva (2008), Weale and Wieladek (2016), and Zhou (forthcoming).

What all the above studies have in common is that they highlight a subset of the identified set as if researchers might be confident that the answer to structural questions of interest falls within that subset. The problem with doing this has been noted by Baumeister and Hamilton (2015, 2018) and Watson (2019). However, the nature of this problem still appears not to be recognized by many applied researchers. For this reason, it is useful to raise these issues again in the context of a particular empirical example in order to explain the nature of the problem and describe the tools available to solve it.

In this paper we do so using for illustration the following question: “What happens to capital flows to emerging markets when U.S. monetary policy becomes more contractionary?” This is an interesting and important question, and the paper by Dahlhaus and Vasishtha (2019) does a nice job of trying to answer it.

When U.S. interest rates go up, as they did for example in May 2013, it could affect emerging market capital flows through several different channels. If the cause of the higher U.S. interest rates was a strengthening of the U.S. economy, the improvement in real investment opportunities in the U.S. could divert some of the capital flows that had been going to emerging markets back to the United States. On the other hand, if the cause of higher U.S. interest rates was a move toward a more contractionary U.S. monetary policy in the absence of any changes in economic fundamentals, that also could induce investors to hold a higher fraction of their assets in U.S. financial instruments. Distinguishing between these channels is one of the goals of Dahlhaus and Vasishtha.

The authors’ approach to identification – that is, the way they propose to separate the effects of different channels like the two described above – uses a combination of sign restrictions and zero restrictions. Their sign restrictions are that a U.S. monetary contraction would increase the 3-year-ahead fed funds futures rate but lower U.S. inflation and output growth. Improving U.S. fundamentals, by contrast, would mean higher interest rates coupled with higher inflation and output growth. Thus the sign restrictions might allow us to distinguish between the two channels. Their additional zero restriction is that a U.S. monetary contraction raises fed funds futures but does not change the current fed funds rate.

Their monetary policy shock is thus in the spirit of a “forward guidance shock” studied by Gürkaynak, Sack, and Swanson (2005) and Swanson (2017).

1 Mechanical details of the algorithm.

We want to focus our discussion on how this identification actually works. To do so we first need to wade into the details underlying the computational algorithm that the authors use. The algorithm is based on that developed for sign-restricted VARs by Uhlig (2005) and Rubio-Ramírez, Waggoner and Zha (2010) as extended to allow for zero restrictions in Baumeister and Benati (2013).

The algorithm is based on a first-order vector autoregression (VAR) using $n = 6$ different variables. These variables are collected in an $(n \times 1)$ vector \mathbf{y}_t consisting of the U.S. fed funds rate, the U.S. 36-month-ahead fed funds futures rate, the U.S. inflation rate, the U.S. industrial production growth rate, the VIX (a measure of U.S. stock-price uncertainty as reflected in the prices of stock options), and the first principal component of a set of capital flows to emerging markets.

Step 1 in this algorithm estimates a reduced-form first-order vector autoregression for these $n = 6$ variables:

$$\underset{(n \times 1)}{\mathbf{y}_t} = \underset{(n \times 1)}{\hat{\mathbf{c}}} + \underset{(n \times n)}{\hat{\mathbf{\Phi}}} \underset{(n \times 1)}{\mathbf{y}_{t-1}} + \underset{(n \times 1)}{\hat{\mathbf{e}}_t} \quad t = 1, 2, \dots, T. \quad (1)$$

The first row of this system is obtained from an ordinary least squares (OLS) regression of the fed funds rate on a constant and one lag of each of the six variables in the VAR. The

first element of $\hat{\mathbf{e}}_t$ is the error one would make in trying to forecast the fed funds rate in month t based on values of the six variables observed at $t - 1$. Associated with these six forecasting regressions is a variance-covariance matrix of the forecast errors,

$$\hat{\mathbf{\Omega}}_{(n \times n)} = T^{-1} \sum_{t=1}^T \hat{\mathbf{e}}_t \hat{\mathbf{e}}_t'.$$

For example, the (1,1) element of $\hat{\mathbf{\Omega}}$ is the average squared error we would make in predicting the fed funds rate on the basis of the OLS regression.

Step 2 in the algorithm is to draw random values $\mathbf{\Omega}^{(m)}$ and $\mathbf{\Phi}^{(m)}$ from the asymptotic distribution of the OLS estimates $\hat{\mathbf{\Omega}}$ and $\hat{\mathbf{\Phi}}$.¹ This step will be repeated many times to generate a number of different draws ($m = 1, 2, \dots, M$).

Step 3 of the algorithm generates a random $(n \times n)$ matrix $\mathbf{Q}^{(m)}$. Each $\mathbf{Q}^{(m)}$ that we draw will be an orthonormal matrix, that is, $\mathbf{Q}^{(m)} \mathbf{Q}^{(m)'} = \mathbf{I}_n$, the $(n \times n)$ identity matrix. The distribution is uniform with respect to a certain measure (known as the Haar measure) over the set of possible orthonormal matrices,² and for this reason the distribution is sometimes thought to be uninformative – more on this below. We then calculate the Cholesky factor $\mathbf{P}^{(m)}$ of the value for $\mathbf{\Omega}^{(m)}$ that was drawn in step 2 (that is, $\mathbf{P}^{(m)}$ is a lower-triangular matrix satisfying $\mathbf{P}^{(m)} \mathbf{P}^{(m)'} = \mathbf{\Omega}^{(m)}$). The proposal is to interpret the observed n reduced-

¹ The asymptotic distribution is the same as the Bayesian posterior distribution that results from using uninformative priors for $\mathbf{\Omega}$ and $\mathbf{\Phi}$. For this reason, this step is sometimes equivalently described as generating draws for $\mathbf{\Omega}^{(m)}$ and $\mathbf{\Phi}^{(m)}$ from the Bayesian posterior distribution obtained using uninformative priors. We will discuss the Bayesian interpretation of the algorithm more below.

² For the case of a VAR with only $n = 2$ variables, the set of possible orthonormal matrices can be summarized as either rotations or reflections with some angle θ . In this case, a draw from the Haar distribution could be obtained by first generating the angle θ from a uniform distribution over the unit circle (that is, all angles are equally likely) and then flipping a fair coin to use this angle either as a reflection or rotation; see Baumeister and Hamilton (2015, p. 1973).

form residuals $\boldsymbol{\varepsilon}_t = \mathbf{y}_t - \mathbf{c} - \boldsymbol{\Phi}\mathbf{y}_{t-1}$ as coming from a set of n structural shocks \mathbf{v}_t according to $\boldsymbol{\varepsilon}_t = \mathbf{P}^{(m)}\mathbf{Q}^{(m)}\mathbf{v}_t$ where the structural shocks each have unit variance and are uncorrelated with each other ($E(\mathbf{v}_t\mathbf{v}_t') = \mathbf{I}_n$). Note that every \mathbf{v}_t we think of in this way is consistent with the observed properties of the reduced-form residuals $\boldsymbol{\varepsilon}_t$, because

$$\begin{aligned} E(\mathbf{P}^{(m)}\mathbf{Q}^{(m)}\mathbf{v}_t\mathbf{v}_t'\mathbf{Q}^{(m)'}\mathbf{P}^{(m)'}) &= \mathbf{P}^{(m)}\mathbf{Q}^{(m)}E(\mathbf{v}_t\mathbf{v}_t')\mathbf{Q}^{(m)'}\mathbf{P}^{(m)'} \\ &= \mathbf{P}^{(m)}\mathbf{Q}^{(m)}\mathbf{Q}^{(m)'}\mathbf{P}^{(m)'} = \mathbf{P}^{(m)}\mathbf{P}^{(m)'} = \boldsymbol{\Omega}^{(m)}. \end{aligned}$$

Thus every proposed structural shock \mathbf{v}_t satisfying $\boldsymbol{\varepsilon}_t = \mathbf{P}^{(m)}\mathbf{Q}^{(m)}\mathbf{v}_t$ also satisfies the condition that $E(\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t') = \boldsymbol{\Omega}^{(m)}$.

Step 4 of the algorithm imposes the sign and zero restrictions. Suppose we label the monetary policy shock as the first element of \mathbf{v}_t . We induce a further rotation of $\mathbf{P}^{(m)}\mathbf{Q}^{(m)}$ such that its (1,1) element (the effect of the monetary shock on the current fed funds rate) is zero, and flip sign of the first column so that the (2,1) element is positive (a monetary contraction should raise the 3-year-ahead fed funds futures rate). We then check whether the (3,1) and (4,1) elements are negative (a monetary contraction should lower both inflation and output). If yes, we keep the first column of $\mathbf{P}^{(m)}\mathbf{Q}^{(m)}$ (denoted $\boldsymbol{\alpha}^{(m)}$) as a plausible effect on impact of a U.S. monetary policy contraction on each of the six variables in the system, and use $[\boldsymbol{\Phi}^{(m)}]^s \boldsymbol{\alpha}^{(m)}$ as a plausible effect on each of the six variables s months after a U.S. monetary contraction. If not, we discard the draw m and generate another.

The structural impulse-response functions for horizon s (shown as the solid lines in Figure 4 in Dahlhaus and Vasishtha) correspond to the median values of retained draws for $[\boldsymbol{\Phi}^{(m)}]^s \boldsymbol{\alpha}^{(m)}$. Their 68% credible sets correspond to the lower 16% and upper 16% values

for the set of retained draws.

2 Implications of the algorithm.

Note that there are two ways in which a random-number generator plays a role in this algorithm. The first is in step 2, in which we drew values for $\Omega^{(m)}$ and $\Phi^{(m)}$ from the asymptotic distribution of the OLS estimates $\hat{\Omega}$ and $\hat{\Phi}$. This step reflects uncertainty in a form that economists are very familiar with, which is sampling uncertainty. We only have a finite number T of observations on \mathbf{y}_t , and because of this we don't know the true values of Ω and Φ . We have an estimate $\hat{\Omega}$, but we understand that the true value might be bigger or smaller than this. To capture the implications of this uncertainty, we generate many possible values $\Omega^{(m)}$, some of which are bigger than $\hat{\Omega}$ and some of which are smaller. We use the distribution of generated $\Omega^{(m)}$ to remind us that the true value might be bigger or smaller than the estimate $\hat{\Omega}$.

A random-number generator is also used in step 3 of the algorithm, which generated a draw for $\mathbf{Q}^{(m)}$ from the Haar distribution on orthonormal matrices. The randomness here came only from the researcher's random-number generator, and has nothing whatever to do with the data or the difficulty we have in estimating objects of interest from a finite sample.

The randomness of $[\Phi^{(m)}]^s \alpha^{(m)}$ (generated draws for the structural impulse-response coefficients) thus comes from a combination of two sources: (1) sampling uncertainty that arises because we only have a finite number of observations and (2) uncertainty that comes purely from a random-number generator used by the researcher.

We can see how much each of these sources of randomness contributes by shutting down the first one altogether.³ Suppose that instead of generating values for $\mathbf{\Omega}^{(m)}$ and $\mathbf{\Phi}^{(m)}$ in step 2 we just fixed $\mathbf{\Omega}^{(m)} = \hat{\mathbf{\Omega}}$ and $\mathbf{\Phi}^{(m)} = \hat{\mathbf{\Phi}}$ for every m . This amounts to proceeding as if we have no sampling uncertainty at all, that is, as if we had an infinite sample of observed data ($T \rightarrow \infty$) so that we could know the true values $\mathbf{\Omega}$ and $\mathbf{\Phi}$ with no sampling uncertainty. Our figure 1 plots the median values for the structural impulse-response function sets for this modification of Dahlhaus and Vasishtha’s algorithm.⁴ On impact a monetary contraction raises the 3-year-ahead fed funds futures rate, lowers inflation and output, and has zero effect on the contemporaneous fed funds rate. All this is guaranteed by the algorithm by construction – a draw would not have been retained unless all of the above were satisfied. Effects generally decay slowly from those impacts – this is a feature of the simple dynamics implied by the OLS estimate of the autoregressive coefficients $\hat{\mathbf{\Phi}}$. The most interesting features might be the last two panels, which show that a U.S. monetary contraction raises U.S. stock price volatility and reduces emerging market capital flows. These conclusions were not imposed by the authors.

The values plotted in Figure 1 are the medians of the set of retained draws. Figure 2 focuses on one magnitude of particular interest – the effect on emerging market capital flows one month after a U.S. monetary policy contraction – and plots the probability distribution of the set of retained draws. The median of this distribution (-0.41) is the value plotted at

³ Similar exercises are reported by Baumeister and Hamilton (2018) and Watson (2019).

⁴ This and the figures below are based on $M = 100,000$ generated draws for $\mathbf{Q}^{(m)}$ of which 21,383 were retained. We thank Tatjana Dahlhaus for generously sharing her data to allow us to perform this exercise.

horizon $s = 1$ in the last panel of Figure 1. But remember that this probability distribution has nothing to do with uncertainty in the data but instead came entirely from the random-number generator used by the researcher. *By construction*, every one of the draws in Figure 2 is perfectly consistent with everything we've observed in the data ($\hat{\Omega}$ and $\hat{\Phi}$) and with all of the identifying restrictions we've imposed. There is no basis whatever from anything we see in the data or anything coming from the identifying restrictions to prefer one of these draws over any other.

And yet the probability distribution in Figure 2 does seem to favor some of these draws over others, seeming to regard values near the median as more likely than others. The only thing that makes them more likely was that the Haar distribution implicitly assumed that they were more likely. Some researchers seem surprised by this observation, since it is commonly believed is that the Haar distribution is uninformative. While it is true that the Haar distribution is uninformative about the angle of rotation associated with the matrix \mathbf{Q} , the object the researcher is interested in is not the angle of rotation of \mathbf{Q} but instead magnitudes like those plotted in Figure 1, namely elements of impulse-response functions. The impulse-response functions are a nonlinear function of the angle of rotation.⁵ A nonlinear transformation of a uniform distribution no longer has a flat distribution, but has more mass for some values than for others.⁶ Whenever researchers use the Haar distribution

⁵ The matrix \mathbf{Q} itself is a nonlinear transformation of the angle of rotation θ . For example, in the case of a rotation matrix for $n = 2$,

$$\mathbf{Q}(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}.$$

⁶ This was one of the main points made by Baumeister and Hamilton (2015).

in an algorithm, they are implicitly assuming that some answers to their question are more plausible than others.

Thus, whenever the set of retained draws is summarized in terms of a median and 68% credibility sets, a researcher is using more than just the information observed in the data and more than the information represented by the identifying sign and zero restrictions. The researcher is also relying on an implicit ranking of possible outcomes. This implicit ranking is an artifact of step 3 of the algorithm described in Section 1. Although many researchers tend to treat step 3 of the algorithm as an incidental mechanical detail, in practice it is often a key factor that led them to draw the conclusions that they did.

3 Possible solutions.

There are four ways that researchers could address this issue: (1) report the identified set; (2) bring in additional weakly identifying information; (3) bring in fully identifying information; or (4) use a combination of the above. We now describe each in turn.

Report the identified set.

If the researcher wants to rely on nothing more than the sign and zero restrictions and what is observed in the data, there is no justification for reporting the median or a 68% interval of the set of retained draws. Instead the researcher should report the full set of all retained draws. For fixed $\hat{\Omega}$ and $\hat{\Phi}$ as here, as the number of generated draws $M \rightarrow \infty$, this corresponds to the set of all the values that are consistent with $\hat{\Omega}$, $\hat{\Phi}$ and the identifying restrictions. This is known as the identified set associated with the fixed $\hat{\Omega}$

and $\hat{\Phi}$. The identified sets for structural impulse-response functions are plotted in Figure 3. The identifying assumptions allow only a single possibility (namely, the number zero) for the effect on impact of the monetary policy shock on the fed funds rate. Thus the identified set for the value at horizon 0 in the first panel in Figure 3 contains only the single value zero. But after one month the effect on the fed funds rate could be positive or negative. The effect on impact on inflation or output of a monetary contraction cannot be positive, but the identified set extends all the way to zero, and after one month the effect could be positive. And the conclusion from the identified set is that the effect of a monetary contraction on VIX or emerging market flows could be essentially anything at near horizons. As the horizon $s \rightarrow \infty$, each identified set shrinks back to a tight interval around zero. This is because $\hat{\Phi}^s \rightarrow \mathbf{0}$ meaning $\hat{\Phi}^s \alpha \rightarrow \mathbf{0}$ for any finite α .

Again we emphasize that this is not at all the usual argument that you can't conclude anything if a 95% confidence interval includes the number zero. We are not talking about a confidence interval in the usual sense here. The usual confidence interval represents uncertainty that would go away if we had an infinite number of observations. We have generated the bounds in Figure 3 assuming we in fact do have an infinite number of observations and that there is no uncertainty at all about the true reduced-form coefficients. The uncertainty instead comes from the fact that, if all we're willing to assume are the stated sign and zero restrictions, then even perfect knowledge of the true reduced-form coefficients would not be enough to determine the magnitude or even the sign of the effect of a monetary policy contraction on emerging market capital flows.

Bring in additional weakly identifying information.

A second alternative is to bring in some additional information not just about the sign but also about the magnitude of certain parameters. This is implicitly what a researcher is doing if the distribution in Figure 2 is treated as if telling us that some elements in the identified set are more likely than others. Indeed, the algorithm in Section 1 can be given a Bayesian interpretation in which the researcher began before seeing the data with some prior beliefs about the relative likelihood of different possible outcomes. The problem with this motivation is that it is far from clear where the information is coming from that enabled us to think some possibilities were more likely to be true than others before we observed the data.

Following this approach correctly would require the researcher to be explicit about what we know from economic theory or other data that gave us a basis for regarding some elements in the identified set as more likely than others. In other words, we need to defend the conclusion that a value like -0.41 in Figure 2 is the most likely value given what we observe in the data and what we expected on the basis of prior information. We may have lots of information, from both economic theory and earlier data sets, that some responses of the U.S. Federal Reserve to inflation and output are more plausible than others. An explicit Bayesian approach allows us to incorporate any prior information like this as well as represent how much confidence we have in that prior information. But this information would be quite unlikely to take the form of a Haar distribution over rotation matrices \mathbf{Q} . Baumeister and Hamilton (2018, 2019) provide a number of examples of how prior information might be

used in practice.

Use fully identifying assumptions.

Another approach is to look for some other variables or information that give us complete identification, that is, information that would enable us, if we knew for certain the values of the reduced-form coefficients, to know for certain the effects of a monetary contraction on emerging market flows. One popular approach developed by Stock and Watson (2012, 2018) and Mertens and Ravn (2014) is to find an instrumental variable that is correlated with a monetary policy shock but uncorrelated with other shocks. Dahlhaus and Vasishtha actually develop such a potential instrument in the form of x_t , which denotes the cumulative changes in the 3-year-ahead fed funds futures on the days in month t when there was a monetary policy announcement. They use this variable as a robustness check, rerunning their sign-restricted VAR in which the second variable in their original system is replaced with x_t . Using x_t in this way results in a system that is still unidentified and that is still subject to the critique raised above.

Another way that a variable like x_t could be used is directly as an instrument for a monetary policy shock.⁷ This produces a framework that is fully identified. Stock and Watson (2012, 2018) and Mertens and Ravn (2014) describe some ways in which this can be done. Noh (2019) and Paul (forthcoming) note that the IV approach is in fact very easy to implement using OLS regressions. Specifically, we add the current value of x_t as an

⁷ For additional discussion of the use of changes in interest rates in a narrow window around FOMC announcements as an instrument for monetary policy shocks, see Nakamura and Steinsson (2018), Hamilton (2018), and Zhang (2019).

additional explanatory variable in each of the regressions in (1):

$$\mathbf{y}_t = \underset{(6 \times 1)}{\tilde{\mathbf{c}}} + \underset{(6 \times 6)}{\tilde{\mathbf{\Phi}}} \underset{(6 \times 1)}{\mathbf{y}_{t-1}} + \underset{(6 \times 1)(1 \times 1)}{\tilde{\boldsymbol{\alpha}}} x_t + \underset{(6 \times 1)}{\tilde{\boldsymbol{\epsilon}}}_t . \quad (2)$$

For example, the first row of (2) is estimated from an OLS regression of the fed funds change at t on a constant, the proxy x_t at date t , and lagged values of the 6 variables in the original system. Noh and Paul showed that if the instrument is valid and relevant, then the estimated value of $\tilde{\mathbf{\Phi}}^s \tilde{\boldsymbol{\alpha}}$ would in an infinite-sized sample be proportional to the true response of each of the 6 variables in \mathbf{y}_{t+s} at date $t+s$ to a monetary shock at date t ; in other words, $\tilde{\mathbf{\Phi}}^s \tilde{\boldsymbol{\alpha}}$ gives a consistent estimate (up to a constant of proportionality) of the magnitude we're interested in. The usual confidence intervals around the estimate $\tilde{\mathbf{\Phi}}^s \tilde{\boldsymbol{\alpha}}$ tell us how uncertain we are about the conclusion given that we've only observed a finite number of observations.

Combining approaches.

It is also possible to combine the best features of the various approaches. For example, we may think that x_t is a reasonable proxy, but we're not completely sure that it's a valid instrument, or may have concerns that it is a weak instrument, in which case we would be back to an unidentified system. Likewise, we may have some other possible zero restrictions that would produce an identified system, but again we may have some doubts about these restrictions. All such approaches can be viewed within a Bayesian context in which we use prior distributions to summarize our uncertainty about the validity of instruments or confidence in zero restrictions. A prior distribution that holds that a certain coefficient is probably close to zero is a strict generalization of an identifying assumption that the

coefficient is exactly equal to zero. Nguyen (2019) demonstrates how to perform inference in a system in which we have doubts about the validity of instruments and doubts about other identifying information. Note that such an approach can be viewed as a strict generalization of both the Rubio-Ramírez, Waggoner and Zha (2010) and Baumeister and Benati (2013) approaches to sign-restricted VARs and the Stock and Watson (2012, 2018) and Mertens and Ravn (2014) approaches to VARs estimated with instrumental variables.

4 Conclusions.

As we noted in the introduction, the method used by Dahlhaus and Vasishtha (2019) is quite well established in the literature. We could have used any of the nearly hundred prominent studies listed there as alternative examples with which to make our points. Although reporting subsets of structural values consistent with sign restrictions has seen widespread acceptance, we think it is useful to remind researchers of the problems with doing so and of the alternatives that avoid these criticisms.

References

Aastveit, Knut Are, Hilde C. Bjørnland, and Leif Anders Thorsrud (2015). “What Drives Oil Prices? Emerging versus Developed Economies,” *Journal of Applied Econometrics* 30: 1013–1028.

Aastveit, Knut Are, Gisle James Natvik, and Sergio Sola (2017). “Economic Uncertainty and the Effectiveness of Monetary Policy,” *Journal of International Money and Finance* 76: 50-67.

Abbate, Angela, Sandra Eickmeier, Wolfgang Lemke, and Massimiliano Marcellino (2016). “The Changing International Transmission of Financial Shocks: Evidence from a Classical Time-Varying FAVAR,” *Journal of Money, Credit and Banking* 48(4): 573–601.

Abbate, Angela, Sandra Eickmeier, and Esteban Prieto (2016). “Financial Shocks and Inflation Dynamics,” working paper, Deutsche Bundesbank.

Abdallah, Chadi S., and William D. Lastrapes (2013). “Evidence on the Relationship between Housing and Consumption in the United States: A State-Level Analysis,” *Journal of Money, Credit and Banking* 45(4): 559–590.

Amir-Ahmadi, Pooyan, Christian Matthes, and Mu-Chun Wang (2016). “Drifts and Volatilities Under Measurement Error: Assessing Monetary Policy Shocks Over the Last Century,” *Quantitative Economics* 7: 591–611.

Amir-Ahmadi, Pooyan, Christian Matthes, and Mu-Chun Wang (2017). “Measurement Errors and Monetary Policy: Then and Now,” *Journal of Economic Dynamics and Control* 79: 66-78.

Antolín-Díaz, Juan, and Juan F. Rubio-Ramírez (2018). “Narrative Sign Restrictions for SVARs,” *American Economic Review* 108: 2802-2829

Anzuini, Alessio, Marco J. Lombardi, and Patrizio Pagano (2013). “The Impact of Monetary Policy Shocks on Commodity Prices,” *International Journal of Central Banking* 9(3): 119-144.

Arias, Jonas E., Dario Caldara, and Juan F. Rubio-Ramírez (2015). “The Systematic Component of Monetary Policy in SVARs: An Agnostic Identification Procedure,” working paper, Emory University.

Arias, Jonas E., Juan F. Rubio-Ramírez, and Daniel F. Waggoner (2018). “Inference Based on SVARs Identified with Sign and Zero Restrictions: Theory and Applications,” *Econometrica* 86: 685-720.

Baumeister, Christiane, and Luca Benati (2013). “Unconventional Monetary Policy and the Great Recession: Estimating the Macroeconomic Effects of a Spread Compression at the Zero Lower Bound,” *International Journal of Central Banking* 9: 165-212.

Baumeister, Christiane, and James D. Hamilton (2015). “Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information,” *Econometrica* 83: 1963-1999.

Baumeister, Christiane, and James D. Hamilton (2018). “Inference in Structural Vector Autoregressions when the Identifying Assumptions Are Not Fully Believed: Re-evaluating the Role of Monetary Policy in Economic Fluctuations,” *Journal of Monetary Economics* 100: 48-65.

Baumeister, Christiane, and James D. Hamilton (2019). “Structural Interpretation of

Vector Autoregressions with Incomplete Information: Revisiting the Role of Oil Supply and Demand Shocks,” *American Economic Review* 109: 1873-1910.

Baumeister, Christiane, and Gert Peersman (2013a). “Time-Varying Effects of Oil Supply Shocks on the US Economy,” *American Economic Journal: Macroeconomics* 5(4): 1-28.

Baumeister, Christiane, and Gert Peersman (2013b). “The Role of Time-Varying Price Elasticities in Accounting for Volatility Changes in the Crude Oil Market,” *Journal of Applied Econometrics* 28: 1087-1109.

Benati, Luca (2015). “The Long-run Phillips Curve: A Structural VAR Investigation,” *Journal of Monetary Economics* 76: 15-28.

Benati, Luca, and Thomas Lubik (2014). “Sales, Inventories and Real Interest Rates: A Century of Stylized Facts,” *Journal of Applied Econometrics* 29(7): 1210–1222.

Belongia, Michael T., and Peter N. Ireland (2016). “The Evolution of U.S. Monetary Policy: 2000–2007,” *Journal of Economic Dynamics and Control* 73: 78-93.

Berg, Tim Oliver (2012). “Did Monetary or Technology Shocks Move Euro Area Stock Prices?” *Empirical Economics* 43(2): 693–722.

Bian, Timothy Yang, and Petro Gete (2015). “What Drives Housing Dynamics in China? A Sign Restrictions VAR Approach,” *Journal of Macroeconomics* 46: 96-112.

Bjørnland, Hilde C., and Jørn Inge Halvorsen (2014). “How Does Monetary Policy Respond to Exchange Rate Movements? New International Evidence,” *Oxford Bulletin of Economics and Statistics* 76(2): 208-232.

Boeckx, Jef, Maarten Dossche, and Gert Peersman (2017). “Effectiveness and Transmis-

sion of the ECB's Balance Sheet Policies," *International Journal of Central Banking* 13(1): 297-333.

Breitenlechner, Max, Johann Scharler, and Friedrich Sindermann (2016). "Banks' External Financing Costs and the Bank Lending Channel: Results from a SVAR Analysis," *Journal of Financial Stability* 26: 228-246.

Buch, Claudia M., Sandra Eickmeier, and Esteban Prieto (2014). "Macroeconomic Factors and Microlevel Bank Behavior," *Journal of Money, Credit and Banking* 46(4): 715-751.

Canova, Fabio, and Gianni De Nicoló (2002). "Monetary Disturbances Matter for Business Fluctuations in the G-7," *Journal of Monetary Economics* 49: 1131-1159.

Canova, Fabio, and Matthias Paustian (2011). "Business Cycle Measurement with Some Theory," *Journal of Monetary Economics* 58: 345-361.

Castelnuovo, Efram, and Paolo Surico (2010). "Monetary Policy, Inflation Expectations and the Price Puzzle," *Economic Journal* 120(549): 1262-1283.

Chadha, Jagjit S., Luisa Corrado, and Qi Sun (2010). "Money and Liquidity Effects: Separating Demand from Supply," *Journal of Economic Dynamics and Control* 34: 1732-1747.

Charnavoki, Valery, and Juan J. Dolado (2014). "The Effects of Global Shocks on Small Commodity-Exporting Economies: Lessons from Canada," *American Economic Journal: Macroeconomics* 6(2): 207-237.

Conti, Antonio M., Stefano Neri, and Andrea Nobili (2017). "Low Inflation and Monetary Policy in the Euro Area," working paper, ECB.

Corsetti, Giancarlo, Luca Dedola, and Sylvain Leduc (2014). “The International Dimension of Productivity and Demand Shocks in the US Economy,” *Journal of the European Economic Association* 12(1): 153-176.

Dahlhaus, Tatjana, and Garima Vasishtha (2019). “Monetary Policy News in the U.S.: Effects on Emerging Market Capital Flows,” working paper, Bank of Canada.

Darracq Paries, Matthieu, and Roberto A. De Santis (2015). “A Non-Standard Monetary Policy Shock: The ECB’s 3-year LTROs and the Shift in Credit Supply,” *Journal of International Money and Finance* 54: 1-34.

Dedola, Luca, and Stefano Neri (2007). “What Does a Technology Shock Do? A VAR Analysis with Model-Based Sign Restrictions,” *Journal of Monetary Economics* 54: 512–49.

Eickmeier, Sandra, and Boris Hofmann (2013). “Monetary Policy, Housing Booms, and Financial (Im)Balances,” *Macroeconomic Dynamics* 17(4): 830–860.

Eickmeier, Sandra, and Tim Ng (2015). “How Do US Credit Supply Shocks Propagate Internationally? A GVAR Approach,” *European Economic Review* 74: 128-145.

Ellis, Colin, Haroon Mumtaz, and Pawel Zabczyk (2014). “What Lies Beneath? A Time-varying FAVAR Model for the UK Transmission Mechanism,” *Economic Journal* 124(576): 668–699.

Enders, Zeno, Gernot J. Müller, and Almuth Scholl (2011). “How do Fiscal and Technology Shocks Affect Real Exchange Rates? New Evidence for the United States,” *Journal of International Economics* 83: 53–69.

Faccini, Renato, Haroon Mumtaz, and Paolo Surico (2016). “International Fiscal Spillovers,”

Journal of International Economics 99: 31-45.

Fadejeva, Ludmila, Martin Feldkircher, and Thomas Reininger (2017). “International Spillovers from Euro Area and US Credit and Demand Shocks: A Focus on Emerging Europe,” *Journal of International Money and Finance* 70: 1-25.

Farrant Katie, and Gert Peersman (2006). “Is the Exchange Rate a Shock Absorber or a Source of Shocks? New Empirical Evidence,” *Journal of Money, Credit and Banking* 38(4): 939–961.

Feldkircher, Martin, and Florian Huber (2016). “The International Transmission of US Shocks—Evidence from Bayesian Global Vector Autoregressions,” *European Economic Review* 81: 167-188.

Fornari, Fabio, and Livio Stracca (2012). “What Does a Financial Shock Do? First International Evidence,” *Economic Policy* 27 (71): 407-445.

Froni, Claudia, Francesco Furlanetto, and Antoine Lepetit (2018). “Labor Supply Factors and Economic Fluctuations,” *International Economic Review* 59: 1491-1510.

Fratzcher, Marcel, Luciana Juvenal, and Lucio Sarno (2010). “Asset Prices, Exchange Rates and the Current Account,” *European Economic Review* 54: 643–658.

Fratzcher, Marcel, and Roland Straub (2013). “Asset Prices, News Shocks, and the Trade Balance,” *Journal of Money, Credit and Banking* 45(7): 1211–1251.

Furlanetto, Francesco, Francesco Ravazzolo, and Samad Sarferaz (2017), “Identification of Financial Factors in Economic Fluctuations,” *Economic Journal* 129: 311-337.

Furlanetto, Francesco, and Ørjan Robstad (2017). “Immigration and the Macroeconomy:

Some New Empirical Evidence,” working paper, Norges Bank.

Fujita, Shigeru (2011). “Dynamics of Worker Flows and Vacancies: Evidence from the Sign Restriction Approach,” *Journal of Applied Econometrics* 26: 89-121.

Gambacorta, Leonardo, Boris Hofmann, and Gert Peersman (2014). “The Effectiveness of Unconventional Monetary Policy at the Zero Lower Bound: A Cross-Country Analysis,” *Journal of Money, Credit and Banking* 46(4): 615–642.

Gambetti, Luca, and Alberto Musso (2017). “Loan Supply Shocks and the Business Cycle,” *Journal of Applied Econometrics* 32(4): 764–782.

Georgiadis, Georgios (2015). “Examining Asymmetries in the Transmission of Monetary Policy in the Euro Area: Evidence from a Mixed Cross-Section Global VAR Model,” *European Economic Review* 75: 195-215.

Glocker, Christian, and Pascal Towbin (2015). “Reserve Requirements as a Macroprudential Instrument – Empirical Evidence from Brazil,” *Journal of Macroeconomics* 44: 158-176.

Güntner, Jochen (2014), “How Do Oil Producers Respond to Oil Demand Shocks?” *Energy Economics* 44: 1-13.

Gupta, Rangan, and Mampho P. Modise (2013). “Does the Source of Oil Price Shocks Matter for South African Stock Returns? A Structural VAR Approach,” *Energy Economics* 40: 825-831.

Gürkaynak, Refet S., Brian Sack, and Eric T. 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: 55-93.

Hamilton, James D. (2018). "The Efficacy of Large-Scale Asset Purchases when the Short-Term Interest Rate is at its Effective Lower Bound," *Brookings Papers on Economic Activity*, Fall 2018: 543-554.

Huber, Florian, and Maria Teresa Punzi (2017). "The Shortage of Safe Assets in the US Investment Portfolio: Some International Evidence," *Journal of International Money and Finance* 74: 318-336.

Hristov, Nikolay, Oliver Hülsewig, and Timo Wollmershäuser (2012). "Loan Supply Shocks During the Financial Crisis: Evidence for the Euro Area," *Journal of International Money and Finance* 31(3): 569-592.

Hofmann, Boris, Gert Peersman, and Roland Straub (2012). "Time Variation in U.S. Wage Dynamics," *Journal of Monetary Economics* 59(8): 769-783.

Jääskelä Jarkko P., and David Jennings (2011). "Monetary Policy and the Exchange Rate: Evaluation of VAR Models," *Journal of International Money and Finance* 30: 1358–1374.

Jarociński, Marek (2010). "Responses to Monetary Policy Shocks in the East and the West of Europe: A Comparison," *Journal of Applied Econometrics* 25(5): 833–868.

Juvenal, Luciana (2011). "Sources of Exchange Rate Fluctuations: Are They Real or Nominal?" *Journal of International Money and Finance* 30: 849–876.

Juvenal, Luciana, and Ivan Petrella (2014). "Speculation in the Oil Market," *Journal of Applied Econometrics* 30(4): 621–649.

Kapetanios, George, Haroon Mumtaz, Ibrahim Stevens, and Konstantinos Theodoridis (2012). “Assessing the Economy-wide Effects of Quantitative Easing,” *Economic Journal* 122(564): F316-F347.

Kilian, Lutz, and Daniel 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(5): 1166-1188.

Kilian, Lutz, and Daniel P. Murphy (2014). “The Role of Inventories and Speculative Trading in the Global Market for Crude Oil,” *Journal of Applied Econometrics* 29: 454-478.

Kilian, Lutz and Xiaoqing Zhou (2019). “Does Drawing Down the U.S. Strategic Petroleum Reserve Help Stabilize Oil Prices?” CEPR Discussion Paper No. DP13849.

Kim, Seong-Hoon, Seongman Moon, and Carlos Velasco (2017). “Delayed Overshooting: Is It an '80s Puzzle?” *Journal of Political Economy* 125: 1570-1598.

Lippi, Francesco, and Andrea Nobili (2012). “Oil and the Macroeconomy: A Quantitative Structural Analysis,” *Journal of the European Economic Association* 10(5): 1059–1083.

Liu, Li, Yudong Wang, Chongfeng Wu, and Wenfeng Wu (2016). “Disentangling the Determinants of Real Oil Prices,” *Energy Economics* 56: 363-373.

Liu, Philip, Haroon Mumtaz, Konstantinos Theodoridis, and Francesco Zanetti (2019). “Changing Macroeconomic Dynamics at the Zero Lower Bound,” *Journal of Business and Economic Statistics* 37: 391-404.

Liu, Philip, Haroon Mumtaz, and Angeliki Theophilopoulou (2014). “The Transmission of International Shocks to the UK. Estimates Based on a Time-Varying Factor Augmented

VAR,” *Journal of International Money and Finance* 46: 1-15.

Luciani, Matteo (2015). “Monetary Policy and the Housing Market: A Structural Factor Analysis,” *Journal of Applied Econometrics* 30(2): 199–218.

Meeks, Roland (2012). “Do Credit Market Shocks Drive Output Fluctuations? Evidence from Corporate Spreads and Defaults,” *Journal of Economic Dynamics and Control* 36: 568–584.

Mertens, Karel, and Morten O. Ravn (2014). “A Reconciliation of SVAR and Narrative Estimates of Tax Multipliers,” *Journal of Monetary Economics* 68: S1-S19.

Michaelis, Henrike, and Sebastian Watzka (2017). “Are There Differences in the Effectiveness of Quantitative Easing at the Zero-Lower-Bound in Japan Over Time?” *Journal of International Money and Finance* 70: 204–233.

Mönch, Emanuel (2012). “Term Structure Surprises: The Predictive Content of Curvature, Level, and Slope,” *Journal of Applied Econometrics* 27: 574–602.

Mountford, Andrew (2005). “Leaning into the Wind: A Structural VAR Investigation of UK Monetary Policy,” *Oxford Bulletin of Economics and Statistics* 67(5): 597-621.

Mountford, Andrew, and Harald Uhlig (2009). “What are the Effects of Fiscal Policy Shocks?” *Journal of Applied Econometrics* 24: 960-992.

Mumtaz, Haroon, and Laura Sunder-Plassmann (2013). “Time-Varying Dynamics of the Real Exchange Rate: An Empirical Analysis,” *Journal of Applied Econometrics* 28: 498–525.

Mumtaz, Haroon, and Francesco Zanetti (2012). “Neutral Technology Shocks and the Dynamics of Labor Input: Results from an Agnostic Identification,” *International Economic*

Review 53(1): 235–254.

Nakamura, Emi, and Jón Steinsson (2018). “High Frequency Identification of Monetary Non-Neutrality: The Information Effect,” *Quarterly Journal of Economics* 133: 1283-1330.

Nam, Deokwoo, and Jian Wang (2016). “Mood Swings and Business Cycles: Evidence from Sign Restrictions,” working paper, Chinese University of Hong Kong.

Nguyen, Lam (2019). “Bayesian Inference in Structural Vector Autoregression with Sign Restrictions and External Instruments,” working paper, UCSD.

Noh, Eul (2019). “Impulse-response Analysis with Proxy Variables,” working paper, SSRN.

Pappa, Evi (2009). “The Effects of Fiscal Shocks on Employment and the Real Wage,” *International Economic Review* 50: 217–244.

Paul, Pascal (forthcoming). “The Time-Varying Effect of Monetary Policy on Asset Prices,” *Review of Economics and Statistics*.

Peersman, Gert (2005). “What Caused the Early Millennium Slowdown? Evidence Based on Vector Autoregressions,” *Journal of Applied Econometrics* 20(2): 185-207.

Peersman, Gert, and Roland Straub (2009). “Technology Shocks and Robust Sign Restrictions in a Euro Area SVAR,” *International Economic Review* 50: 727–750.

Riggi, Marianna, and Fabrizio Venditti (2015). “The Time Varying Effect of Oil Price Shocks on Euro-area Exports,” *Journal of Economic Dynamics and Control* 59: 75-94.

Rubio-Ramírez, Juan F., Daniel F. Waggoner, and Tao Zha (2010). “Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference,” *Review of Economic*

Studies 77(2): 665-696.

Sá, Filipa, Pascal Towbin, and Tomasz Wieladek (2014). “Capital Inflows, Financial Structure and Housing Booms,” *Journal of the European Economic Association* 12(2): 522-546.

Sá, Filipa, and Tomasz Wieladek (2015). “Capital Inflows and the U.S. Housing Boom,” *Journal of Money, Credit and Banking* 47(s1): 221–256.

Schenkelberg, Heike and Sebastian Watzka (2013). “Real Effects of Quantitative Easing at the Zero Lower Bound: Structural VAR-based Evidence from Japan,” *Journal of International Money and Finance* 33: 327-357.

Scholl, Almuth, and Harald Uhlig (2008). “New Evidence on the Puzzles: Results from Agnostic Identification on Monetary Policy and Exchange Rates,” *Journal of International Economics* 76: 1–13.

Stock, James H., and Mark W. Watson (2012). “Disentangling the Channels of the 2007-09 Recession,” *Brookings Papers on Economic Activity*, 43(1):81–156.

Stock, James H., and Mark W. Watson (2018). “Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments,” *Economic Journal* 128: 917-948.

Swanson, Eric T (2017). “Measuring the Effects of Federal Reserve Forward Guidance and Asset Purchases on Financial Markets,” National Bureau of Economic Research Working Paper 23311.

Uhlig, Harald (2005). “What are the Effects of Monetary Policy on Output? Results

from an Agnostic Identification Procedure,” *Journal of Monetary Economics* 52: 381-419.

Van Robays, Ine (2016). “Macroeconomic Uncertainty and Oil Price Volatility,” *Oxford Bulletin of Economics and Statistics* 78(5): 671–693.

Vargas-Silva, Carlos (2008). “Monetary Policy and the US Housing Market: A VAR Analysis Imposing Sign Restrictions,” *Journal of Macroeconomics* 30: 977–990.

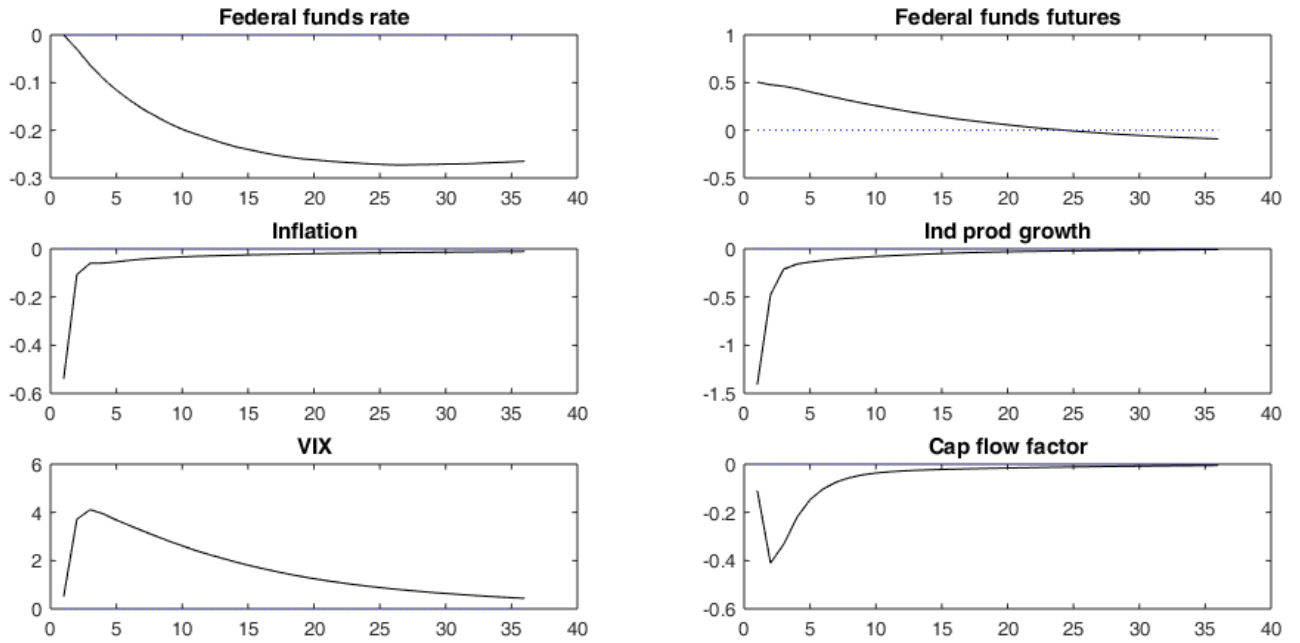
Watson, Mark (2019). “Comment on ‘On the Empirical (Ir)Relevance of the Zero Lower Bound’ by D. Debortoli, J. Gali, and L. Gambetti”, *NBER Macroeconomics Annual*.

Weale, Martin, and Tomasz Wieladek (2016). “What Are the Macroeconomic Effects of Asset Purchases?” *Journal of Monetary Economics* 79: 81-93.

Zhang, Xu (2019). “Disentangling the Information Effects in the Federal Reserve’s Monetary Policy Announcements,” working paper, UCSD.

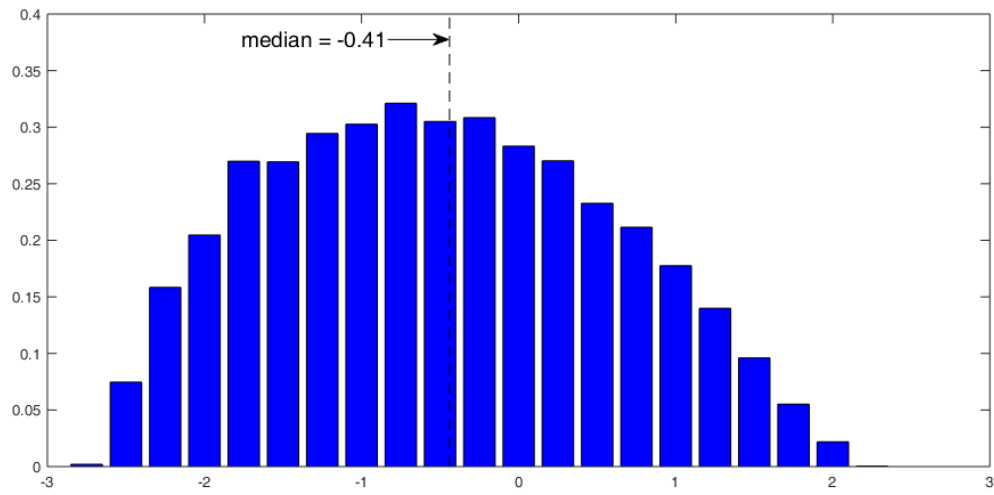
Zhou, Xiaoqing (forthcoming). “Refining the Workhorse Oil Market Model,” *Journal of Applied Econometrics*.

Figure 1. Median structural impulse-response functions when there is no parameter uncertainty.



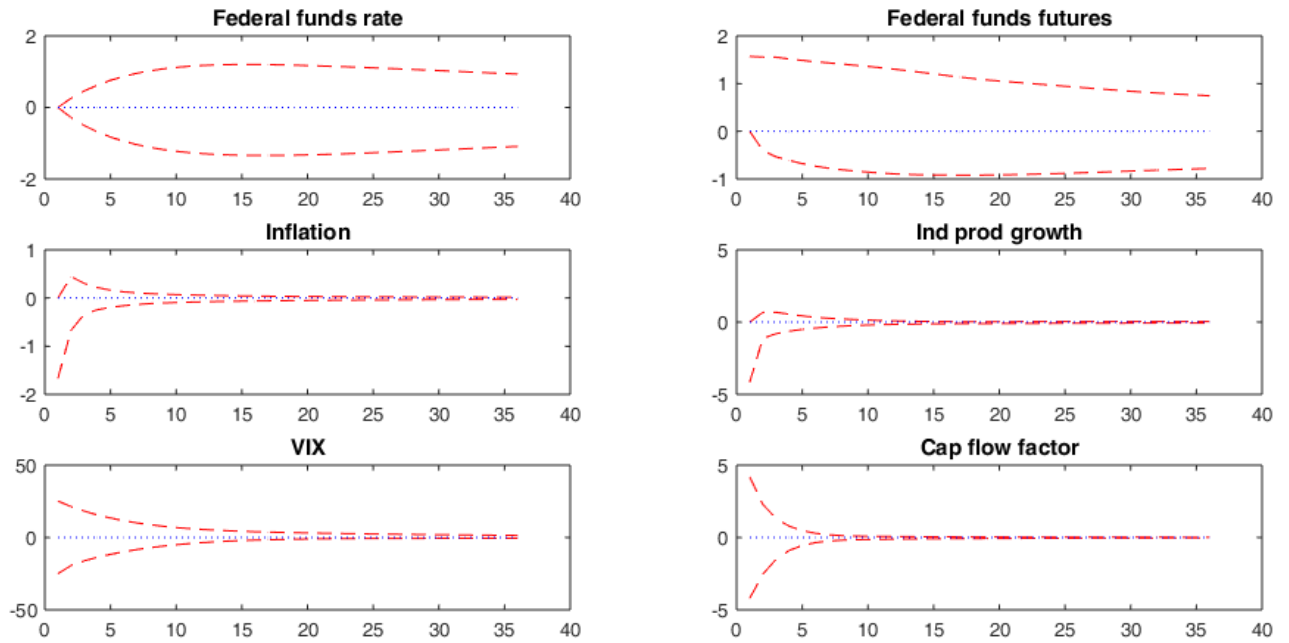
Notes to Figure 1. For each horizon s months the i th panel in the figure plots the median value of the i th element of $[\Phi^{(m)}]^s \alpha^{(m)}$ generated by the algorithm in Section 1 modified by setting $\Omega^{(m)} = \hat{\Omega}$ and $\Phi^{(m)} = \hat{\Phi}$ for every m .

Figure 2. Distribution of draws for effect of monetary policy shock on capital flows one period later when there is no parameter uncertainty.



Notes to Figure 2. The figure plots the probability distribution of the set of draws for the horizon $s=1$ term of the last panel in Figure 1.

Figure 3. Identified sets for structural impulse-reponse functions when there is no parameter uncertainty.



Notes to Figure 3. The figure plots the upper and lower bounds for each s of draws from the algorithm used to generate Figure 1.