Marek Jarociński and Peter Karadi

Deconstructing Monetary Policy Surprises - The Role of Information Shocks

Online Appendix

BAYESIAN ESTIMATION

This section explains how we estimate the VAR in (1) and reports some prior sensitivity analysis.

B1.   The baseline prior

The VAR in (1) in matrix notation is

(B1)
$$\begin{pmatrix} M & Y \end{pmatrix} = X \begin{pmatrix} 0 & B \end{pmatrix} + \begin{pmatrix} U^m & U^y \end{pmatrix}.$$

where $M = (m_1, ..., m_T)'$, $Y = (y_1, ..., y_T)'$, X is a matrix that collects the right-hand-side variables, with a typical row $x_t' = (m_{t-1}'\, y_{t-1}' \,... \, m_{t-P}'\, y_{t-P}'\, 1)$, $B = \left(B_{YM}^1, B_{YY}^1, ..., B_{YM}^P, B_{YY}^P, c_y\right)'$, $U^m = (u_1^m, ..., u_T^m)'$, and $U^y = \left(u_1^y, ..., u_T^y\right)'$. Let m^o denote the vector collecting the *observed* values in M and m^* denote a vector collecting the *missing* values in M.

The prior about B and Σ is independent normal-inverted Wishart, $p(B, \Sigma) = p(B)p(\Sigma)$, where

(B2)
$$p(\Sigma|\underline{S}, \underline{v}) = \mathcal{IW}\left(\underline{S}, \underline{v}\right) \propto |\Sigma|^{-\underline{v}/2} \exp\left(-\frac{1}{2}\operatorname{tr}\underline{S}\Sigma^{-1}\right),$$

(B3)
$$p(\operatorname{vec} B|\underline{B}, \underline{Q}) = \mathcal{N}(\operatorname{vec}\underline{B}, \underline{Q}) \propto \exp\left(-\frac{1}{2}\operatorname{vec}(B - \underline{B})'\underline{Q}^{-1}\operatorname{vec}(B - \underline{B})\right),$$

$\mathcal{IW}$ denotes the Inverted Wishart distribution and $\mathcal{N}$ denotes the normal distribution.

In $\underline{B}$ the coefficient of the first own lag of each variable is 1 and the remaining entries are zero. $\underline{Q}$ is a diagonal matrix implying that the standard deviation of lag p of variable j in equation i is $\lambda_1^{-1}\sigma_i/\sigma_j p^{-\lambda_2}$. Following Litterman (1986) we take $\lambda_1 = 5$, $\lambda_2 = 1$ and σ_i (σ_j) is the standard error in the autoregression of order P of variable i (j). We set $\underline{v} = N + 2$ and $\underline{S}$ is a diagonal matrix with σ_i^2, $i = 1, ..., N_m + N_y$ on the diagonal.

To handle the missing values in M we cast model (1) in the state-space form. The prior about m^*, $p(m^*|B, \Sigma)$ is implicit from model (1) (**?**, Ch.8). We assume that the initial values are $m_{-P+1} = ... = m_0 = 0$.

We use a Gibbs sampler to compute the posterior. The Gibbs sampler consists of drawing in turn Σ, B and m^* from their conditional posteriors until the sampler converges.

B2.   The conditional posteriors

The conditional posteriors are as follows.

- The conditional posterior of Σ:

(B4)
$$p(\Sigma|Y, M, B) = \mathcal{IW}\left(\overline{S}, \overline{v}\right)$$

where

(B5) $\quad \overline{S} = \left(\begin{pmatrix} M & Y \end{pmatrix} - X \begin{pmatrix} 0 & B \end{pmatrix}\right)' \left(\begin{pmatrix} M & Y \end{pmatrix} - X \begin{pmatrix} 0 & B \end{pmatrix}\right) + \underline{S},$

(B6) $\quad \overline{v} = T + \underline{v}.$

- The conditional posterior of B:

(B7)
$$p(\operatorname{vec} B|Y, M, \Sigma) = \mathcal{N}\left(\overline{B}, \overline{Q}\right)$$

where

(B8) $\quad \overline{Q} = \left(\underline{Q}^{-1} + \Sigma_{YY.1}^{-1} \otimes X'X\right)^{-1},$

(B9) $\quad \operatorname{vec} \overline{B} = \overline{Q}\left(\underline{Q}^{-1} \operatorname{vec} \underline{B} + \left(\Sigma_{YY.1}^{-1} \otimes X'\right) \operatorname{vec}\left(Y + M\Sigma_{MM}^{-1}\Sigma_{MY}\right)\right)$

and we use the notation $\Sigma = \begin{pmatrix} \Sigma_{MM} & \Sigma_{MY} \\ \Sigma_{YM} & \Sigma_{YY} \end{pmatrix}$ and $\Sigma_{YY.1} = \Sigma_{YY} - \Sigma_{YM}\Sigma_{MM}^{-1}\Sigma_{MY}$.

- The conditional posterior of m^* is given by the simulation smoother. We use the simulation smoother of Durbin and Koopman (2002) implemented as explained in Jarociński (2015).

B3. Derivation of the conditional posteriors

The conditional posteriors of Σ and m^* are standard.

To obtain the conditional posterior of B we write down the density of Y, M conditional on the parameters B and Σ

(B10) $\quad p(Y, M|B, \Sigma) \propto |\Sigma|^{-T/2}$
$$\exp\left(-\frac{1}{2}\operatorname{tr}\left(\begin{pmatrix} M & Y \end{pmatrix} - X \begin{pmatrix} 0 & B \end{pmatrix}\right)' \left(\begin{pmatrix} M & Y \end{pmatrix} - X \begin{pmatrix} 0 & B \end{pmatrix}\right)\Sigma^{-1}\right).$$

and decompose it as follows:

(B11)
$$p(Y, M|B, \Sigma) = p(Y|M, B, \Sigma)p(M|B, \Sigma)$$

where

(B12)
$$p(M|B, \Sigma) = p(M|\Sigma_{MM}) \propto |\Sigma_{MM}|^{-T/2}\exp\left(-\frac{1}{2}\operatorname{tr} M'M\Sigma_{MM}^{-1}\right)$$

2

and

(B13) $p(Y|M, B, \Sigma) \propto |\Sigma_{YY.1}|^{-T/2}$

$$\exp\left(-\frac{1}{2}\operatorname{tr}\left(Y - XB + M\Sigma_{MM}^{-1}\Sigma_{MY}\right)'\left(Y - XB + M\Sigma_{MM}^{-1}\Sigma_{MY}\right)\Sigma_{YY.1}^{-1}\right)$$

with $\Sigma_{YY.1} = \Sigma_{YY} - \Sigma_{YM}\Sigma_{MM}^{-1}\Sigma_{MY}$. See e.g. Bauwens, Lubrano and Richard (1999) Section A.2.3.

We notice that the only terms in the posterior that involve B are $p(Y|M, B, \Sigma)p(B)$. We multiply them out and collect the terms involving B in the standard way.

<center>B4.   Prior sensitivity analysis</center>

In this section we use a more general prior and report the marginal data densities for alternative hyperparameter values. We follow the analysis and the notation of Del Negro and Schorfheide (2011). Specifically, we add to the prior a 'sums-of-coefficients dummy observation prior' with weight λ_4 and a 'co-persistence dummy observation prior' with weight λ_5.
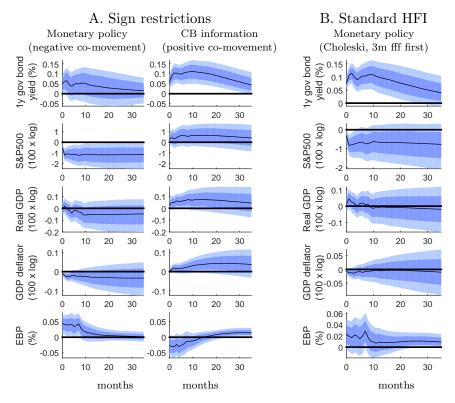
We adapt these priors to our setup of a VAR with zero restrictions and an independent (and hence, non-conjugate) normal-inverted Wishart prior. That is, we write down the new $\underline{Q}$ and $\underline{B}$ that reflect both the Litterman (1986) prior and the dummy observation priors. More in detail, we specify the dummy observations Y^d, X^d that correspond to the 'sums-of-coefficients' prior in the equations for y_t and to the 'co-persistence' prior. Let $\underline{\Sigma}$ denote the error variance in the dummy observations sample. We assume that $\underline{\Sigma}$ equals the prior expectation of the error variance in the estimation sample, i.e. $\underline{\Sigma} = E(\Sigma) = \underline{S}$, where $\underline{S}$ is a diagonal matrix described in section B.B1. Let $\underline{\Sigma}_{YY}$ be the part of the variance matrix that corresponds to the equations for y_t. Let $\underline{Q}^L$ and $\underline{B}^L$ denote the variance and mean of the Litterman's prior described in section B.B1. Combining the Litterman's normal prior with the likelihood of the dummy observations we obtain a normal prior for B with the variance and mean given by

(B14) $$\underline{Q} = \left((\underline{Q}^L)^{-1} + \underline{\Sigma}_{YY}^{-1} \otimes X^{d\prime}X^d\right)^{-1},$$

(B15) $$\operatorname{vec}\underline{B} = \underline{Q}\left((\underline{Q}^L)^{-1}\operatorname{vec}\underline{B}^L + (\underline{\Sigma}_{YY}^{-1} \otimes X^{d\prime})\operatorname{vec}Y^d\right).$$

Figure B1 shows that the impulse responses change modestly when we add to the prior the dummy observation priors with weights $\lambda_4 = \lambda_5 = 1$. These are the weights used e.g. in Sims and Zha (1998) and these weights also approximately maximize the marginal data density in our application. Comparing Figure B1 with Figure 2 we see two main differences. One is that the responses of output and prices to the central bank information shock become more persistent with the dummy observation priors. Another difference is that the responses of output and prices to the monetary policy shocks become less negative. This happens both in our sign restriction identification (panel A) and in the standard high frequency identification (panel B). In our identification the responses of these variables remain marginally significant, but they become basically zero in the standard high frequency identification. Hence, under the Sims and Zha (1998) prior our sign restrictions become also qualitatively, and not only quantitatively important.

Table B1 reports the marginal data density for several specifications of the prior. We compute the marginal data density for a small grid of values for $\lambda_1, \lambda_2, \lambda_4$ and λ_5. We use the modified

<center>3</center>

FIGURE B1. IMPULSE RESPONSES OF THE LOW FREQUENCY VARIABLES y_t TO MONETARY POLICY AND CENTRAL BANK INFORMATION SHOCKS, MODEL WITH DUMMY OBSERVATION PRIORS WITH SIMS AND ZHA (1998) HYPERPARAMETERS. MEDIAN (LINE), PERCENTILES 16-84 (DARKER BAND), PERCENTILES 5-95 (LIGHTER BAND).

harmonic mean estimator based on a chain with 100 thousand draws after discarding the first 5000, using the observed-data likelihood.[35] The marginal likelihood in our VAR is quite sensitive to λ_1 and λ_2, but rather insensitive to λ_4 and λ_5 for the values that we have tried. The first lesson from Table B1 is that the Sims and Zha (1998) specification is the approximate local maximum. The marginal data density goes down when we either tighten or loosen the hyperparameters. The second lesson is that the Sims and Zha (1998) prior is considerably, though not overwhelmingly, preferred to the baseline specification. The approximately 7 log points difference is substantial, though not huge by the standards of the marginal data densities. All in all, we have decided to keep the specification without the dummy observation priors as the baseline in the main text, as in this specification the results of the standard high frequency identification (in Panel B of Figure 2) are closer to the literature, which often uses frequentist econometrics. In this way we focus this paper on the conceptual advantages of our sign restriction identification and not on the performance of the standard high frequency identification in the Bayesian framework. We leave the latter topic for future research.

---

[35]In the applications with missing data, like ours, one can also use the complete-data likelihood based on the draws of the missing data, but Chan and Grant (2015) argue strongly for using the observed-data likelihood.

	λ_1	λ_2	λ_4	λ_5	$\ln p_\lambda(Y)$
tighter λ_4, λ_5	5	1	3	3	-880.1
tighter λ_2	5	2	1	1	-885.9
tighter λ_1	7	1	1	1	-895.0
Sims and Zha (1998)	5	1	1	1	-879.7
looser λ_1	3	1	1	1	-899.4
looser λ_2	5	0.5	1	1	-944.4
looser λ_4, λ_5 (baseline)	5	1	0	0	-887.2

ADDITIONAL RESULTS FOR THE US

C1. *Relaxing the restrictions on the dynamics of m_t*

In this subsection we show that our results are robust to relaxing the restrictions on the dynamics of m_t in the VAR. The unrestricted VAR is

$$
(C1) \qquad \begin{pmatrix} m_t \\ y_t \end{pmatrix} = \sum_{p=1}^{P} \begin{pmatrix} B_{MM}^p & B_{MY}^p \\ B_{YM}^p & B_{YY}^p \end{pmatrix} \begin{pmatrix} m_{t-p} \\ y_{t-p} \end{pmatrix} + \begin{pmatrix} c_M \\ c_Y \end{pmatrix} + \begin{pmatrix} u_t^m \\ u_t^y \end{pmatrix}.
$$

We estimate this VAR on the sample without the missing values in m_t, i.e. starting in February 1990. Furthermore, we replace the missing observation in September 2001 with zero. In this way we can estimate a completely standard VAR. Panel A of Figure C1 reports the resulting impulse responses. Panel B reports the impulse responses obtained with the restricted VAR given in equation (1) on the sample starting in February 1990. We can see that the impulse responses in both panels are extremely similar. We conclude that relaxing the zero restrictions in the VAR hardly affects the impulse responses.

An additional lesson from Figure C1 is that starting the sample in 1990 does not change the conclusions. We can see that the impulse responses in this figure are quite similar to the impulse responses in Figure 2.

C2. *Results on other subsamples*

Figure C1 showed that the findings hardly change when we start the sample in February 1990 instead of February 1984. Figure C2 shows that the findings continue to be similar when we estimate the VAR on a sample that starts in February 1984 but ends on December 2008, i.e. before the interest rates hit the zero lower bound (ZLB) in January 2009 (Panel A). Furthermore, the findings continue to be similar when we omit the high-frequency surprises before February 1994 (Panel B). The motivation to omit these surprises is that the Fed did not issue a press release about FOMC decisions until February 1994, so the earlier surprises might be coming from a different regime. Finally, the findings continue to be similar when we start the sample in July 1979, as in the related work by Gertler and Karadi (2015) (Panel C).

5

A. Unrestricted VAR given in equation (C1)

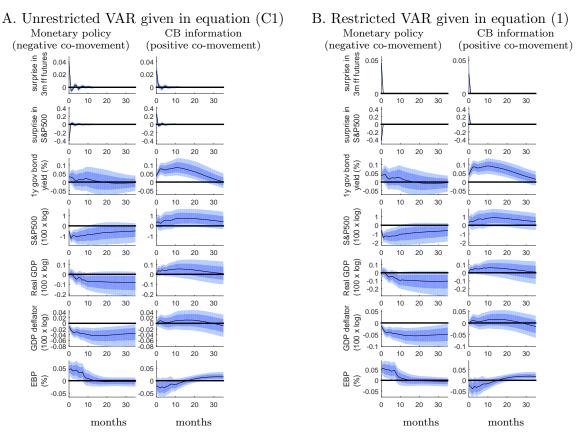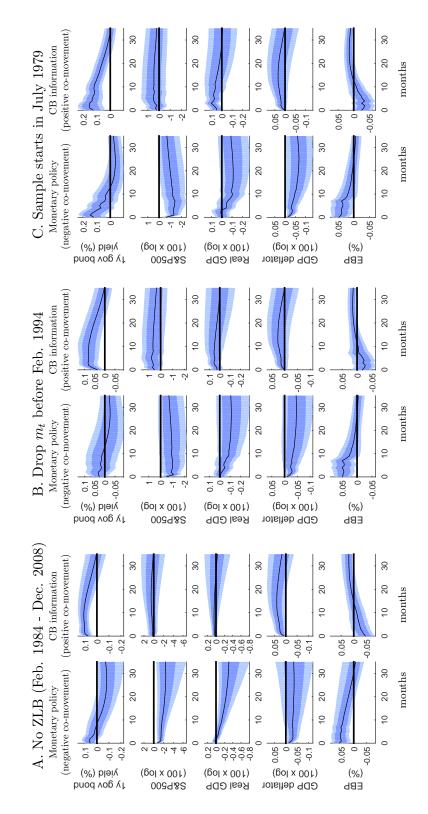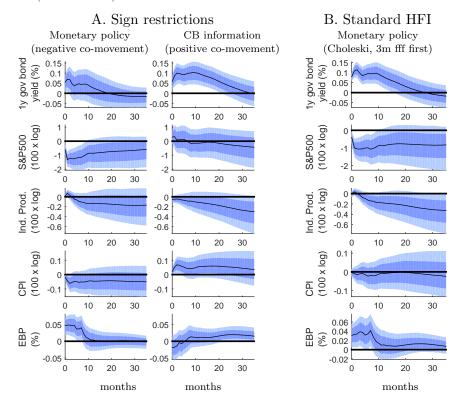B. Restricted VAR given in equation (1)

FIGURE C2. IMPULSE RESPONSES OF THE LOW FREQUENCY VARIABLES y_t TO MONETARY POLICY AND CENTRAL BANK INFORMATION SHOCKS: RESULTS FOR SUBSAMPLES. MEDIAN (LINE), PERCENTILES 16-84 (DARKER BAND), PERCENTILES 5-95 (LIGHTER BAND).

7

C3. *Results with Industrial Production and CPI*

Figure C3 shows that when we replace the real GDP and GDP deflator with the industrial production and the consumer price index (CPI), the standard high-frequency identification yields no response of consumer prices, while these prices do respond in our identification scheme.

FIGURE C3. IMPULSE RESPONSES OF THE LOW FREQUENCY VARIABLES y_t TO MONETARY POLICY AND CENTRAL BANK INFORMATION SHOCKS, MODEL WITH INDUSTRIAL PRODUCTION AND CONSUMER PRICE INDEX. MEDIAN (LINE), PERCENTILES 16-84 (DARKER BAND), PERCENTILES 5-95 (LIGHTER BAND).

A. Sign restrictions

Monetary policy (negative co-movement)

CB information (positive co-movement)

B. Standard HFI

Monetary policy (Choleski, 3m fff first)

C4. VAR with factors of high-frequency surprises

This section shows the robustness of our results to alternative measures of surprises.

FACTORS OF HIGH-FREQUENCY SURPRISES. — We start by showing that the proportion and sizes of 'wrong-signed' responses of stock prices to monetary policy surprises remain similar when we use alternative measures of surprises.

As an alternative measure of the interest rate surprises we compute the 'policy indicator' constructed as in Nakamura and Steinsson (2018) (who build on Gürkaynak, Sack and Swanson, 2005a). Namely, this is the first principal component of the surprises in fed funds futures and eurodollar futures with one year or less to expiration. Five indicators enter into it: the current-month fed funds future, the 3-month fed funds future, and the eurodollar futures at the horizons of two, three and four quarters. The advantage of the policy indicator is that it captures even more of the forward guidance. The disadvantage is that it relies on the eurodollar futures which are not as liquid as the federal funds futures.

As an alternative measure of the stock price surprises we take the first principal component of the surprises in the S&P500, Nasdaq Composite and Wilshire 5000. Nasdaq Composite is based on about 4000 stocks skewed towards the technology sector, and Wilshire 5000 is based on 7000 stocks of essentially all publicly listed companies headquartered in the US. All three indices are market capitalization-weighted. Our dataset has many missing values for Nasdaq and Wilshire, so we use the alternating least squares (ALS) algorithm that simultaneously estimates the missing values while computing principal components.

Table C1 reports the correlations between the 3-month fed funds futures surprises, S&P500 surprises and the two alternative measures of surprises just discussed. The correlation between the surprises in the 3-month fed funds futures and the policy index is 0.89. The correlation between S&P500 and the first principal component of the three stock indices is higher, 0.96. The correlations between interest surprises and stock price surprises are between -0.4 and -0.5.
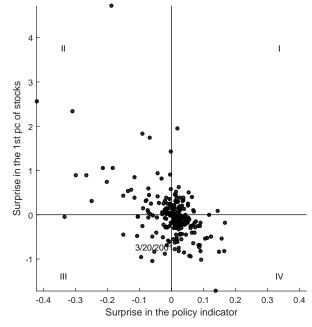
TABLE C1—CORRELATIONS BETWEEN SURPRISES

	3-m fff	SP500	Policy indicator	1st p.c. of stocks
3-m fff	1.00			
SP500	-0.46	1.00		
Policy indicator	0.89	-0.53	1.00	
1st p.c. of stocks	-0.40	0.96	-0.47	1.00

Figure C4 shows that when we use the alternative measures of surprises, the lessons on the 'wrong-signed' responses of stock prices to interest rates hold. Still, in 33% of the cases the co-movement between interest rates and stock price surprises is positive. This confirms the lessons from Figure 1.

IMPULSE RESPONSES. — Now we use the factors extracted from multiple interest rate and stock market surprises as m_t in the VAR. Figure C5 shows that using factors changes very little in the

Note: Each dot represents one FOMC announcement.

impulse responses. The main difference is that the monthly S&P500 index now responds positively to the central bank information shock.

C5. *Robust error bands of Giacomini and Kitagawa (2015)*

This section shows that the impulse responses to the two shocks we identify continue to be very different from each other irrespective of the prior on the rotation matrices Q. We make this point using the 'multiple priors' approach of Giacomini and Kitagawa (2015).

The prior on Q might be important, because the sign restrictions in Table 1 provide only a set identification, not a sharp identification. That is, for every nonsingular variance matrix Σ there is a continuum of rotation matrices Q that are consistent with the sign restrictions. Since the sample carries no information about Q, the weights on different values of Q are determined by the prior. As most of the literature, we use the uniform prior on the space of rotation matrices, conditionally on satisfying the sign restrictions (Rubio-Ramirez, Waggoner and Zha, 2010). How much could the impulse responses change if we used some other, non-uniform prior on Q?

To answer this question we compute the 'robust' uncertainty bounds following Giacomini and Kitagawa (2015). In this approach, the *posterior mean bounds* delineate the range of the posterior means of the impulse responses across all possible priors on Q that satisfy the sign restrictions. The *X% robustified region* is a range of values of the impulse responses that has the posterior probability of at least X% under every possible prior on Q that satisfies the sign restrictions.

Figure C6 reports the robust bounds for the impulse responses of all variables y_t at all horizons. The bounds are wider and include zero more often than the bounds in Figure 2, but the different
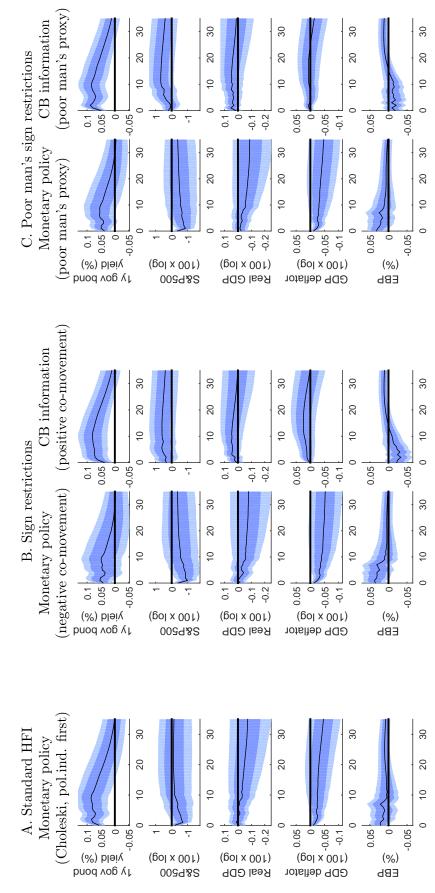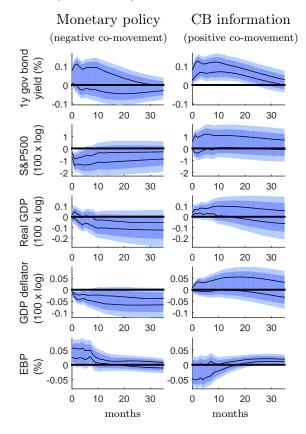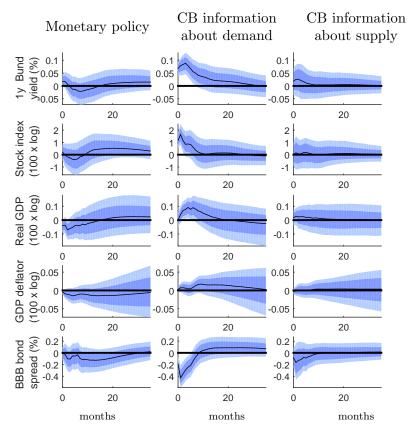
FIGURE C5. IMPULSE RESPONSES OF THE LOW FREQUENCY VARIABLES y_t TO ONE STANDARD DEVIATION SHOCKS, VAR WITH FACTORS OF SURPRISES. MEDIAN (LINE), PERCENTILES 16-84 (DARKER BAND), PERCENTILES 5-95 (LIGHTER BAND). MONTHS ON THE HORIZONTAL AXIS.

11

nature of the monetary policy and central bank information shocks remains clear. Furthermore, let us make two comments related to the width of the bounds. First, the robust bounds are conservative because they account for the 'worst-case' prior on Q for each variable, shock and horizon separately. Any single prior on Q will produce narrower bands. Second, there are many ways to refine the sign restriction identification by postulating further reasonable restrictions on the impulse responses. Our point in this paper is that the simple sign restriction we propose is enough to separate two shocks of very different nature.

Figure D1. Impulse responses of the low frequency variables y_t to one standard deviation shocks. Median (line), percentiles 16-84 (darker band), percentiles 5-95 (lighter band). VAR with central bank information about supply.

Surprises and proxies for Fed's private information

In this section we study the relation between a popular proxy for the private information available to the FOMC members and the central bank information shocks we identify. We find mixed results.

Empirical proxies for the FOMC private information used in the literature are based on the differences between the Fed staff forecasts and private forecasts. For every scheduled FOMC meeting, the Fed staff prepares nowcasts and forecasts of the price level and economic activity. These forecasts do not directly influence private forecasts, because they are made public only with a 5 year delay. However, they are made available to the FOMC members, who can take them into account when setting the course of policy and formulating official communication. The staff forecasts have been shown to have superior forecasting ability relative to private forecasts (Romer and Romer, 2000). The difference between the staff forecasts and forecasts of private forecasters, therefore, is a popular proxy for the private information of the FOMC. Controlling for private information using

these proxies has been shown to influence predictions about the effects of monetary policy shocks (Gertler and Karadi, 2015; Campbell et al., 2016).

It is far from clear, however, how much of the FOMC private information is actually revealed through a policy change and the accompanying communication. FOMC decision makers might not share the views of the staff about the economy, and even if they do their communication might not be detailed enough to explain all the assumptions behind their choices. Therefore, instead of using such proxies, we use market-price reactions to the announcements to learn about the information content of the FOMC statements in our baseline regressions. Changes in asset prices provide more first-hand signal about the extent of new information in the statement as assessed by market participants (and not just by economic forecasters), who can be expected to have key influence on market prices that drive economic fundamentals. Still, it is worthwhile to assess how well our measures line up with private information proxies.

TABLE E1—SURPRISES AND PROXIES FOR FED PRIVATE INFORMATION

Variables	(1) Surprise in 3m fff	(2) Monetary policy shock	(3) CB information shock
π_t	0.00203	0.00209	0.000288
	(0.330)	(0.383)	(0.0660)
π_{t+1}	0.00623	0.00163	0.00497
	(0.474)	(0.201)	(0.776)
π_{t+2}	-0.00799	-0.00514	-0.00363
	(-0.835)	(-0.849)	(-0.717)
dy_t	0.0181***	0.0183***	-0.00141
	(2.893)	(3.119)	(-0.388)
dy_{t+1}	0.0140	0.000733	0.0143***
	(1.379)	(0.0886)	(3.078)
dy_{t+2}	-0.00758	-0.00220	-0.00671
	(-0.891)	(-0.341)	(-1.643)
u_t	-0.0279	-0.0256	-0.00629
	(-0.630)	(-0.796)	(-0.296)
Observations	180	180	180
R-squared	0.117	0.116	0.070

Robust t-statistics in parentheses
*** p<0.01, ** p<0.05, * p<0.1

To this end, we regress the surprises in the 3-month fed funds futures and our two identified shocks on proxies for the FOMC private information. The variables are at the monthly frequency. As measures of the two shocks we take the posterior medians of the respective shocks' contributions to the surprises in the 3-month fed funds futures. The proxy for the FOMC private information is standard in the literature. In particular, we link the staff forecasts on scheduled FOMC meetings

with the last preceding forecasts surveyed by the Blue Chip Economic Indicators. We use the current, and the one- and two-quarters ahead GDP deflator $(\pi_t, \pi_{t+1}, \pi_{t+2})$ and real GDP growth $(dy_t, dy_{t+1}, dy_{t+2})$ forecasts and the current month unemployment forecasts (u_t). We take a simple difference between the staff and private forecasts for each variable. The regression results are shown in Table E1.

The results are mixed. We find that private information about the one-quarter-ahead real GDP growth influences the central bank information shocks significantly. At the same time, we do not find that private information about prices or the unemployment rate would influence the same shock; and we also find that private information about the current-quarter real GDP growth influences our monetary policy shock.

<center>High-frequency euro area data</center>

We use high-frequency data on euro area asset prices to build a dataset of high-frequency asset price responses to the ECB policy announcements, analogous to the Gürkaynak, Sack and Swanson (2005a) dataset for the US. We take the high-frequency asset price data from the Thomson Reuters Tick History database. Our dataset has two kinds of assets: interest rate swaps and stock prices.

Stock prices. For the stock prices it is straightforward to obtain high-frequency data, since stocks are traded in centralized markets. The stock index we use is Euro Stoxx 50. The Thomson Reuters includes its price multiple times a second.

Interest rate swaps. In the euro area we use the interest rate swaps instead of the futures, as the swap market is more liquid and has a longer history. We use the Overnight Indexed Swaps (OIS) based on the Eonia rate. In this swap contract the parties exchange the variable, overnight Eonia rate for the fixed swap rate. We focus on the 3-month swap.

Measuring the Eonia OIS rate is more difficult than measuring stock prices, because these swaps are traded in over-the-counter markets. We do not observe the prices. Thomson Reuters only provides the quotes posted by individual traders. The quotes consist of a bid rate and an ask rate, and the trades are concluded over the phone. The database includes bid and ask quotes with time stamps (at the millisecond level) and with the identity of the posting institution. Some quotes are outliers that cannot reasonably reflect actual trades (e.g. they differ from the other quotes at that time by orders of magnitude). To clean the data from the outliers, for each day, we exclude the lowest and highest 1 percents of bid and ask quotes. In some instances, we eliminate further outliers if they are very far from the outstanding quotes (sometimes 5-6 standard deviations away) making it unreasonable to assume that any trade was conducted at the quoted price.

We measure the market price as the average of the highest bid and lowest ask prices out of the most recent five quotes made by distinct institutions. Furthermore, we disregard quotes posted more than 15 minutes ago, even if this reduces the number of available quotes below 5. In the instances when the highest bid price is higher than the lowest ask price we go for the second-highest and second-lowest or third-highest and third-lowest if necessary. Our choices are informed by our aim to obtain an accurate and timely proxy for market valuation. Choosing the five latest quotes balances timeliness with accuracy: if after a market news 5 institutions modified their quotes, we would like our measure to reflect the change, even if some still outstanding quotes (possibly posted before the news) suggest different valuations. We disregard quotes older than 15 minutes altogether, because quotes can not be directly traded on. They are indicative of the valuation of the posting institution only when they were made, and can lose their actuality over time. The 15 minutes limit

<center>15</center>

guarantees that our baseline surprise measure, which reads the asset price 20 minutes after the monetary policy news, does not include quotes made before the news.
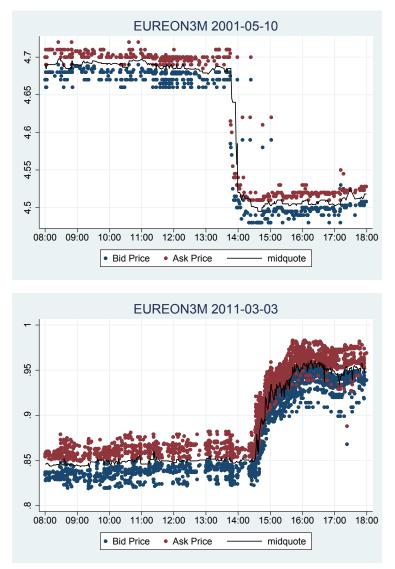
Figure F1 shows two examples illustrating how we process the data on quotes. Each quote is represented by a pair of dots: a blue dot, showing the bid rate, and a red dot, showing the ask rate. The outliers are already removed, as they would distort the scale of the picture. The black line shows the midquote, which is our measure of the market rate. The first panel presents the market for the 3-month Eonia OIS (EUREON3M) on May 10th, 2001. On that day the ECB announced a 25 basis point cut in its policy rates. The press release was issued at 13:45. We can see that around 13:45 the quotes drop by about 20 basis points. The midquote we compute drops with

16

the quotes. The second panel shows the data for March 3rd, 2011. The activity in the market is higher in 2011 than in 2001, as witnessed by a much larger number of quotes posted. On this particular day the ECB Governing Council decided to keep the policy rates unchanged. This was anticipated, so the press release at 13:45 did not contain any surprises. However, during the press conference that started at 14:30 and lasted about an hour, the ECB President Jean-Claude Trichet delivered a hawkish message. He highlighted the upwards risks to inflation coming from an increase in commodity prices, and concerns about second-round effects (i.e. the price increases fuelling wage increases). By the end of the press conference the 3-month Eonia OIS was about 10 basis points higher, reflecting expectations of future interest rate increases.

TABLE G1—CALIBRATED MODEL PARAMETERS

		Households
β	0.990	Discount rate
h	0.815	Habit parameter
χ	3.411	Relative utility weight of labor
φ	0.276	Inverse Frisch elasticity of labor supply
$\bar{S}_h/S$	0.500	Relative steady state direct HH holding of debt
$\varrho_{k,x}$	0.974	Rate of geometric decline of a corporate bond with duration x
		Financial Intermediaries
θ	0.343	Fraction of capital that can be diverted
ω	0.0019	Start-up fund for the entering bankers
σ	0.972	Survival rate of the bankers
		Intermediate good firms
α	0.330	Capital share
δ	0.025	Depreciation rate
		Capital Producing Firms
η_i	1.728	Inverse elasticity of net investment to the price of capital
		Retail Firms
ϵ	4.167	Elasticity of substitution
		Government
$\frac{G}{Y}$	0.200	Steady state proportion of government expenditures
κ_π	1.500	Inflation coefficient in the Taylor rule
κ_x	-0.125	Markup coefficient in the Taylor rule