Has the Transmission of US Monetary Policy Changed Since 2022?

Philip Barrett* Josef Platzer†

January 31, 2025

Abstract

Activity and inflation responded slowly to the Federal Reserve's rate hikes in 2022. Was this because the transmission of monetary policy had changed? Or did other shocks offset tighter policy? We use pre-pandemic data to estimate a VAR with monetary policy shocks identified from high-frequency data, and use it as a filter to back out the sequence of structural monetary policy shocks consistent with data since 2022. We compare these implied shocks to the actual shocks and find the difference statistically significant during February-July 2022. These differences imply that monetary transmission was around 25 percent weaker than normal. Our method accounts for other shocks; allowing them to change to match the post-COVID covariance of the data produces similar results. We provide suggestive evidence that weaker monetary policy transmission was the result of a steeper Phillips curve.

JEL-Codes: C32, E43, E52
Keywords: Monetary Policy, Semi-structural Identification, VAR, Filtering

*International Monetary Fund, Washington, D.C. 20431, USA, email: pbarrett@imf.org
†International Monetary Fund, Washington, D.C. 20431, USA, email: jplatzer@imf.org, corresponding author

The views expressed herein are those of the authors and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

1 Introduction

In March 2022, as inflation reached its highest level since the mid-1980s, the Federal Reserve made the first of a series of interest rate hikes. In the next 19 months, the federal funds rate increased from zero to over five percent. By any standards, this was a large and fast monetary tightening. In the four preceding tightening cycles increases in the federal funds rate had been much smaller and slower, taking an average of 21 months to reach peaks typically 2.5 percentage points above their starting levels.[1]

By July 2022, it was clear that this tightening cycle would last some time. By then, many forecasters were predicting that interest rates would increase further and remain elevated until at least late-2023. This policy stance was widely expected to reduce output below its long-run trend and pull inflation back down towards target, in line with standard views of the transmission of monetary policy (see Figure 1).

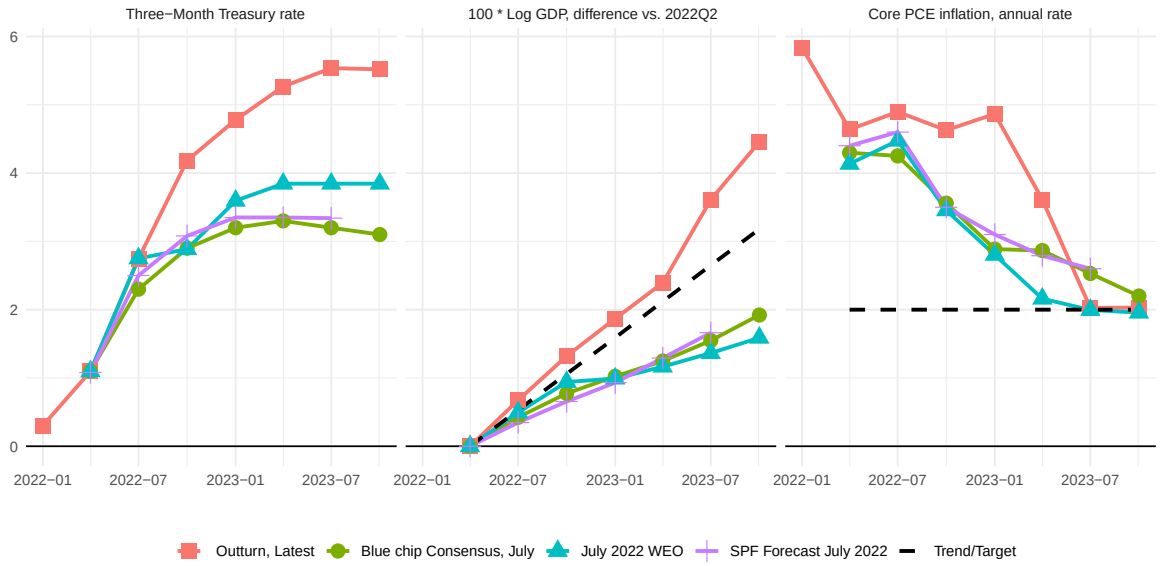


Figure 1: Outturns versus forecasts made in July 2022

Figure shows outturns for interest rates, output, and the price level versus forecasts made in July 2022. Forecasts for change in log GDP abstract from revisions by calculating relative to contemporaneous data for forecasts. The trend for log GDP is at annualized rate of 2.1, the 2000-2019 average. For core PCE the target is 2 percent, the Federal Reserve's inflation target. Abbreviations: *WEO*, IMF's World Economic Outlook; *SPF*, Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters.

However, outturns largely confounded these predictions. Output grew rapidly throughout 2022 and 2023, at more than double the rate of even the most optimistic forecasters. This growth surprise occurred even though interest rates ended up being much higher than expected, something which would normally be expected to have dampened real activity fur-

[1] Cycles starting in February 1994, June 1999, June 2004, and December 2015

ther. And despite the surge in output, inflation has since returned to levels close to the Federal Reserve's two percent target.

The exceedingly strong performance of the real economy and the apparently unconnected disinflation during such an aggressive tightening cycle raises an obvious question: was the transmission of monetary policy somehow different during this episode? And if so, how? This paper seeks to answer these questions.

The key challenge answering these questions is in disentangling the impact of other macroeconomic shocks (or combinations of them) from any potential change in the transmission of monetary policy. For example, the mix of an unanticipated fiscal expansion and a positive supply shock could produce higher-than-expected output and interest rates and falling inflation, even if monetary policy had been transmitted exactly as expected.

Given this challenge, we develop a method to answer this question which accounts for the impact of other shocks but without having to identify them explicitly. Our basic approach proceeds in three steps. We first construct a series of monetary policy shocks from high-frequency intra-day financial market data in the spirit of Jarociński and Karadi (2020) and use a pre-COVID sample to estimate a factor-augmented vector autoregression (FAVAR) which includes the monetary policy shocks as data.[2] The Cholesky decomposition of the FAVAR's reduced form shock with the monetary shocks ordered first then has a semi-structural interpretation where the monetary policy shocks are identified but non-monetary ones are not. The purpose of this step is to produce a data generating process for macroeconomic time series which embodies relatively standard views about the transmission of monetary policy in the pre-COVID period.[3]

In the second step, we invert the FAVAR to infer the most likely sequence of monetary shocks given the macro data alone. This can be thought of as a thought experiment where we pretend that we do not see the monetary shocks starting in 2021.[4] We then generate the distribution of the high-frequency shocks since 2021 under the null hypothesis that the mapping from monetary policy to the macro data has not changed. A strength of our method is that it explicitly takes account of the changing composition of the other shocks. Although the inferred monetary shocks depend directly on the estimated transmission of monetary policy, they depend on other shocks only indirectly via the covariance of the reduced form residuals. Put differently, even if some unusual combination of other shocks were driving outcomes in the post-tightening period, this would not affect the inferred monetary shocks at

[2]Aside from the monetary policy shocks, the baseline VAR also includes data on the fed funds rate, industrial production, private consumption, inflation, and the excess bond premium as well as the first five principal components of a large set of other macroeconomic time series.

[3]As Paul (2020) shows, including the instrument directly in the VAR is identical to the more standard external instruments approach pioneered by Stock and Watson (2012) and Mertens and Ravn (2013), up to re-scaling by a constant. We discuss the relationship of our method to using external instruments further in Section 2.

[4]Despite the title of our paper, we start in 2021 to give something of a lead in to the 2022 tightening cycle.

all, so long as those other shocks were consistent with the estimated covariance of the data. This feature suggests a natural extension: to relax this assumption simply by replacing the estimated reduced form covariance matrix with another one, such as one estimated on the post-2021 data. This allows a calculation of the distribution of monetary shocks consistent with the data and unchanged transmission, even if the impact of other shocks has also changed. Moreover, this can be done without specifying exactly how other shocks have changed, just their impact on the covariance of the aggregate data.

In the third step, we compare these distributions to the observed high-frequency monetary policy shocks since 2021. If the null hypothesis holds—that monetary policy transmission has not changed—-then the realized and inferred shocks should line up. We propose maximum likelihood tests to assess the strength of the evidence that the transmission of monetary policy has changed, as well as modifications to these tests which permit a quantitative assessment of how much weaker monetary policy might have been.

We present results for two types of monetary shocks within the same estimation framework, one a "pure" shock to the stance of monetary policy and the other a Central Bank information shock. The former corresponds to truly tighter monetary policy whereas the latter occurs when the Federal Reserve communicates its insights about other future shocks to other agents in the economy. We find that between February and July 2022, the measured high-frequency monetary stance shocks were persistently and statistically significantly higher than those implied by the filter, both on average and in (some) individual months. That is, although the Federal Reserve was repeatedly surprising the markets with unexpectedly tight monetary policy, the macro data were responding as if it were not. This suggests that monetary policy was less effective than usual during 2022. We also show that the realized Central Bank information shocks were larger than those implied by the filter throughout 2021; the data responded as if the Fed was communicating much more negatively than it did.

We investigate the results for the pure monetary shock further, providing evidence on the magnitude of the change in the strength of monetary policy. We show that in our baseline specification, one can reject at the 5 percent confidence level a null that during February-July 2022 monetary policy was up to 25 percent less effective. This means that for every three interest rate hikes that the Fed did during this period, they would have needed to do approximately one more to deliver the same degree of tightening as one would have expected pre-COVID. When we allow for the variance of other shocks to change to replicate the covariance of the 2021-2023 data, we find that this conclusion is a little weaker – holding with the same confidence only in March and April 2022. Conversely, including policymakers' speeches in our shock series strengthens this finding, with the same effect holding at the one percent level.

We also analyze the source of the uncertainty arising in our approach, since the quantifi-

cation of this uncertainty is essential to formulating the hypothesis tests we apply. We find that uncertainty over the model parameters accounts for a slim minority of the total. Instead, uncertainty due to the colinearity of shocks — that is the extent to which monetary policy and other shocks have similar impact on the data and are thus hard to disentangle—are slightly more important.

To try to understand why transmission might have changed, we calibrate a well-known New Keynesian macro model (Smets and Wouters (2007)) to match our impulse responses. We then replace key structural parameters with alternate values to produce a series of alternate data generating processes which we can use as candidate null hypotheses in our test of the post-2021 data. This allows us to answer the question: if one had a structural model which captured the pre-COVID transmission of monetary policy, what parameter would need to be changed (and by how much) in order to give monetary transmission consistent with the data when monetary policy was tightening. Of course, this is not equivalent to full re-estimation of a structural model, so we consider this evidence to be suggestive rather than conclusive. Nevertheless, we find that steeper price and, to a lesser extent, wage Phillips curve are able to explain the post-2022 data.

Related Literature Our work contributes to three prior areas of research.

First, there is an emerging literature on the monetary policy response to the COVID-19 pandemic and its aftermath (English et al. (2024), Ball et al. (2022)). Much of this work focuses on key ingredients of the monetary transmission mechanism, rather than the more general notion of transmission that we look at. For example, Stedman and Pollard (2023) highlight that services sector labor markets have been particularly tight since the pandemic. These sectors are less sensitive to changes in interest rates, potentially dampening the monetary policy impact. Similarly, Cohen (2023) argues that resilience in the labor market is connected to severe labor shortages in the post-pandemic era, leading employers to hold on to workers and hire less-skilled workers. D'Amico and King (2023) argue that the lags to monetary policy might have been shorter in the post-pandemic tightening cycle because of an increased role of the expectations channel of monetary policy. Glover and Oliyide (2024) note that private-lending spreads have been unusually low in this period, implying less restrictive monetary policy. Bauer et al. (2024) find that public perception of the Fed's monetary policy response to inflation has shifted over the post-pandemic period. Such a shift can explain disinflation with relatively low output and unemployment costs. Fosso et al. (2023) ask a similar question as our paper, but employ a different methodology. They use a structural VAR, allowing for time-varying parameters, and don't find evidence for a change in monetary policy transmission, but that offsetting shocks can explain why unemployment has stayed low in the post-pandemic period, despite rapid policy tightening. These studies of the post-COVID period also tie into a broader literature on the state-

dependent effects of monetary policy. Although there is a considerable evidence in favor of impact of the state-dependent impact of monetary policy, there is not clear agreement on what the states are and why they matter.[5] Analyses which specify *a priori* that monetary transmission is affected by a particular state risk ignoring other sources of variation in transmission. By investigating whether monetary policy has operated differently at specific periods, rather than the more general issue of in specific states, we sidestep the preliminary question of what that state should be.

Second, a methodological literature on using high-frequency monetary policy shocks as identified shocks in vector autoregressive (VAR) models (Stock and Watson (2012), Mertens and Ravn (2013), Gertler and Karadi (2015)). These papers study the dynamic impact of identified monetary policy shocks on the economy. While we use a similar approach to estimate a VAR, our ultimate application is different. Once we estimated the VAR, we invert it and use it as a filter to back out the most likely set of shocks over a test period. We then construct a statistical test on whether the mapping of monetary policy to macro data has changed or not. Similar to Paul (2020), our pursuit of an extra step not easily represented in the external instruments framework motivates our direct inclusion of shocks in the VAR.

Third, we contribute to the literature that uses high-frequency changes in asset prices around specific events to construct identified monetary policy shocks (Kuttner (2001), Gürkaynak et al. (2005) and many others). Initially the literature was centered on Federal Open Market Committee (FOMC) announcements. However, there are only a limited number of scheduled FOMC meetings (typically eight per year), and the size of the asset price changes around the events is small. And so a recent literature (Swanson (2023), Swanson and Jayawickrema (2023)) extends the approach to include post-FOMC-meeting press conferences, releases of FOMC meeting minutes, and speeches by the Chair and Vice Chair of the Federal Reserve. They find that these events are important sources of variation for asset prices. In an extension, we expand the list of events further by including speeches by all FOMC Board members.

The paper is organized as follows: in Section 2 we describe our method; in Section 3 we discuss the construction of our monetary policy shocks; in Section 4 we describe the estimated FAVAR and validate our method on the pre-COVID sample; in Section 5 we present our main results, comparing the filter-implied shocks with those measured in the high frequency data; in Section 6 we we discuss alternative interpretations of our results; and in Section 7 we conclude.

[5]For example, Tenreyro and Thwaites (2016) show that monetary policy is more powerful in expansions than recessions, something that Bernstein (2021) rationalizes with a model with occasionally-binding liquidity constraints. Ascari and Haber (2022) connect to the literature on endogenous price-setting and show that monetary policy is more effective when either shocks are large or when trend inflation is low. And Berger et al. (2021) and Eichenbaum et al. (2022) highlight that the presence of substantial debt in fixed-rate, prepayable mortgages can lead to state-dependent effects of monetary policy.

2 Methodology

2.1 Estimating a Semi-structural Vector Autoregression

Let y_t be a vector of N macro variables and v_t is an i.i.d. vector of proxies for K different monetary policy shocks identified from other data. We describe how to create such a series in the next section but here we take this as given. We collect monthly data on y_t and v_t since 1990 (more detail on the data we use in section 4.1) We call the period from 1990-2019 the *estimation period* and 2021-2023 the *testing period*, with lengths $N_{est} = 336$ and $N_{test} = 36$ respectively.

We assume that the data is generated by a VAR of the form:

$$\begin{bmatrix} v_t \\ y_t \end{bmatrix} = \sum_{j=1}^{J} A_j \begin{bmatrix} v_{t-j} \\ y_{t-j} \end{bmatrix} + \epsilon_t, \qquad \epsilon_t \sim N(0, \Omega) \tag{1}$$

and impose that

$$A_j = \begin{bmatrix} 0_{K\times K} & 0_{K\times N} \\ 0_{N\times K} & B_j \end{bmatrix}$$

The top line of zeros follows from the assumption that v_t is i.i.d. The zeros in the other rows simply impose that y_t depend only on lagged shocks through lagged y_t, not directly (i.e. that y_t has a reduced form VAR, not VARMA, representation). We estimate this specification on the estimation period.

As is usual, we assume that the reduced form errors are a linear function of the true, structural shocks, so that:

$$\epsilon_t = D\delta_t \qquad\qquad Cov\ \delta_t \sim N(0, I_{N+K}) \tag{2}$$

Where $DD' = \Omega$. Because v_t are already externally-identified monetary shocks, any decomposition of Ω where the non-diagonal elements of the top K rows are all zero will have monetary shocks in the first two columns.[6] If the shocks are orthogonal, a Cholesky decomposition will satisfy this condition.

Then we write:

$$D = \begin{bmatrix} I_K & 0 \\ D_m & D_x \end{bmatrix}$$

Then, the D_m columns thus describe the direct impact of monetary policy shocks on the macroeconomy in the estimation period. We also will find it helpful to use similar notation

[6] Strictly speaking, the first K entries in δ_t are the *innovations* to v_t. But since the shocks are i.i.d. this distinction is not meaningful.

to divide ϵ_t:

$$\epsilon_t = \begin{bmatrix} \epsilon_t^v \\ \epsilon_t^y \end{bmatrix}$$

And for δ_t, we use the superscripts ^{m} and ^{x} to denote the monetary and non-monetary shocks:

$$\delta_t = \begin{bmatrix} \delta_t^m \\ \delta_t^x \end{bmatrix}$$

Later, we will consider the case where we pretend that we only see the reduced form errors for the macro variables, ϵ_t^y. Then, we can rewrite equation (2) as:

$$D_y \delta_t = \epsilon_t^y \tag{3}$$

where $D_y = [D_m \; D_x]$. Note that because ϵ_t^y is of length N but δ_t is of length $N+K$, equation (3) does not have a unique solution for δ_t. Instead, a set of structural shocks are consistent with any given observation for ϵ_t^y. In the next section we characterize the distribution of δ_t within that set.

The relationship to external instruments: A common alternative approach is to use v_t as an external instrument (Stock and Watson (2012), Mertens and Ravn (2013), and Olea et al. (2021)). That is, one could rewrite equation (1) purely in terms of the y_t,

$$y_t = \sum_{j=1}^{J} B_j y_{t-j} + \nu_t, \qquad \nu_t \sim N(0, \Xi) \tag{4}$$

And if the elements of v_t are strongly correlated with the monetary shocks (instrument strength) but plausibly uncorrelated with the other structural shocks (the exclusion restriction), one can use the v_t as an external instrument for the monetary shocks. The result is a decomposition of Ξ where the first two columns can be interpreted as the responses to a policy shocks.

Paul (2020) shows – subject to some technical conditions satisfied in our application – that the resulting impulse responses are identical to directly including the v_t series in the VAR, up to a scalar constant. That is, the two approaches differ only in their unit of measurement rather than anything more fundamental.[7] Given that the approaches are in this sense equivalent (something which we check in our estimation later), we prefer the direct inclusion of the shock in the VAR because it makes our next step – recasting the VAR as

[7] A yet further equivalent formulation is one where the data generating process for y_t includes an explicit contemporaneous relationship, in which case $\tilde{B}_0 y_t = \sum_{j=1}^{J} \tilde{B}_j y_{t-j} + \tilde{\nu}_t$. This is identical to equation (4) if $\tilde{B}_j = \tilde{B}_0 B_j$, $\tilde{\nu}_t \sim N(0, \tilde{B}_0 \Xi \tilde{B}_0')$ and so the same equivalence between external instruments and inclusion of the external shocks in the VAR still holds.

a filter – much easier. A similar extension motivated the derivation of the results in Paul (2020). He estimates a VAR with time-varying parameters, which is much easier when the identified shocks are directly included rather than used as instruments.

2.2 Recasting the VAR as a Filter

We now pretend that we do not observe the monetary policy shock on the testing period. The purpose of this assumption is to use the macro data y_t alone to infer what the monetary shocks would have been under the null hypothesis that the transmission of monetary policy was unchanged.

Conditional on the estimates $\hat{\theta} = \{\hat{B}_1, \ldots, \hat{B}_J, \hat{\Omega}\}$, we can compute the fitted reduced form errors $\hat{\epsilon}_t^y$ associated with the macro data y_t alone by:

$$\hat{\epsilon}_t^y = y_t - \sum_{j=1}^{J} \hat{B}_j y_{t-j}$$

As noted above, we cannot solve for a unique δ_t given ϵ_t^y. Instead we calculate the most likely structural shock given $\hat{\theta}$, via maximum likelihood.

Because $\delta_t \sim N(0, I_{N+K})$, the maximum likelihood function just reduces to least squares. Thus, we can compute the most likely structural shocks by solving the following linear-quadratic program period-by-period:

$$\hat{\delta}_{t|\hat{\theta}} = \arg \min_{\delta \in \mathbb{R}^{N+K}} \delta'\delta \qquad \text{s.t.} \quad \hat{D}_y \delta = \hat{\epsilon}_t^y$$

where $\hat{D}_y$ is the sample equivalent of D_y, i.e. the bottom N rows of the Cholesky decomposition of $\hat{\Omega}$. This has solution:

$$\hat{\delta}_{t|\hat{\theta}} = \hat{D}_y'(\hat{\Omega}_y)^{-1}\hat{\epsilon}_t^y \tag{5}$$

Where $\hat{\Omega}_y = \mathbb{E}\hat{\epsilon}_t^y \, \hat{\epsilon}_t^{y\prime} = \hat{D}_y \hat{D}_y'$ is the estimated variance covariance matrix of the reduced form residuals for the data, ϵ_t^y. Because we are only interested in the monetary shock, we will find it helpful to focus only on the top K rows of this estimator:

$$\hat{\delta}^m_{t|\hat{\theta}} = \hat{D}_m'(\hat{\Omega}_y)^{-1}\hat{\epsilon}_t^y$$

It is straightforward to show that this is unbiased, conditional on unbiasedness of $\hat{\theta}$. That is:

$$\mathbb{E}\hat{\theta} = \theta \qquad \Rightarrow \qquad \mathbb{E}\hat{\delta}^m_{t|\hat{\theta}} = \delta_t^m$$

We can also compute the mean square error of the estimated structural shocks:

$$\mathbb{E}_t\left(\hat{\delta}^m_{t|\hat{\theta}} - \delta^m_t\right)\left(\hat{\delta}^m_{t|\hat{\theta}} - \delta^m_t\right)' = I_K - \hat{D}'_m\hat{\Omega}_y^{-1}\hat{D}_m \tag{6}$$

This is just a Kalman filter with the unobserved structural shocks as a hidden state and the reduced form shocks as observed signals. The only differences from the usual set-up are that there is no persistence in the state variable and that uncertainty over the state comes not from noise but from incomplete observation (in that the dimension of $\hat{\epsilon}^y_t$ is less than $\hat{\delta}_t$).[8]

Given that δ_t itself is normal and that $\hat{\delta}^m_{t|\hat{\theta}}$ is an unbiased estimator for δ^m_t conditional on the data and the model, then the distribution of the true shocks conditional on the data and the model parameters, which we denote $\delta^m_{t|t,\hat{\theta}}$, is given by a multivariate normal:

$$\delta^m_{t|t,\hat{\theta}} \sim \mathcal{N}\left(\hat{D}'_m(\hat{\Omega}_y)^{-1}\hat{\epsilon}^y_t,\ I_K - \hat{D}'_m\hat{\Omega}_y^{-1}\hat{D}_m\right) \tag{7}$$

We will find it helpful to denote the corresponding normal density by $f_{t,|\hat{\theta}}(\delta)$.

An important property of this filter is that the distribution for the inferred monetary shocks are a function only of D_m and the reduced form error covariance $\hat{\Omega}_y$. One would get exactly the same results for *any* other $\tilde{D}_x$, satisfying:

$$\tilde{D}_x\tilde{D}'_x = D_xD'_x$$

In other words, the filtered monetary shocks are wholly invariant to any other interpretation of the non-monetary shocks, so long as they are consistent with the data covariance. Moreover, if we want to know how any other candidate *distribution* of the non-monetary shocks affects our results we need only specify its impact on the data via the reduced form residual covariance matrix, $\hat{\Omega}_y$. We will use this fact later when investigating the robustness of our results to the changing variance of non-monetary shocks.

2.3 Parameter Uncertainty

The preceding section derives the maximum likelihood estimator and the mean square error for the structural shocks given the data and a model. However, the true model is unknown. The distribution in equation (7) does not account for this. And so to account for model uncertainty, we integrate over the distribution of the parameters via bootstrap.

Specifically, we use the point estimates $\hat{\theta} = \{\hat{B}_1, \ldots, \hat{B}_J, \hat{\Omega}\}$ to simulate N_{sim} draws of the estimation sample $\{(v^i_1, y^i_1), \ldots, (v^i_{N_{est}}, y^i_{N_{est}})\}_{i=1}^{N_{sim}}$ from the data generating process in equa-

[8] A related literature is that on conditional forecasting in VARs, initiated by Waggoner and Zha (1999), where one seeks to forecast a subset of variables in a VAR by imposing conditioning assumptions on the other variables. Our approach maps to this framework if one interprets the inferred shocks as the subset of variables to be forecast and the observed macro data as the conditioning assumptions.

tion (1). On each sample, we estimate the corresponding parameters $\hat{\theta}^i = \{\hat{B}_1^i, \ldots, \hat{B}_J^i, \hat{\Omega}^i\}$. We then compute the distribution for $\delta_{t|t}$, the structural shocks conditional only on the data. This is a mixture of normal distributions with density given by:

$$f_t(\delta^m) = \frac{1}{N_{sim}} \sum_{i=1}^{N_{sim}} f_{t|\hat{\theta}^k}(\delta^m) \tag{8}$$

This density represents the distribution of shocks under the null that the relationship between the data and the shocks remains unchanged, taking into account uncertainty around that relationship (i.e. model uncertainty). As such, we can use it to assess hypothesis tests of this null.

This approach has two further benefits. First, we can use the simulated parameter draws to compute booststrapped confidence intervals for impulse responses. Second, in each testing period t we can compare the distributions $f_{t,\hat{\theta}}(\cdot)$ and $f_t(\cdot)$ to understand the sources of uncertainty over our estimates of the structural shocks. The former captures the irreducible uncertainty coming from the colinearity of the structural shocks. The latter adds the effect of model uncertainty, which is mitigated in longer samples. We return to this point later in Section 5.4

3 Constructing Monetary Policy Shocks from High-Frequency Financial Data

3.1 Approach

We identify monetary shocks using the high-frequency identification approach pioneered by papers including Kuttner (2001) and Gürkaynak et al. (2005). The idea is to look at the change in the price of securities around a narrow time window of relevant monetary policy events, such as FOMC meetings. The change in price should capture the pure surprise component of the news announcement uncorrelated with other shocks, and therefore is a good candidate for an identified monetary policy shock. The choice of a narrow window aims to eliminate the possibility that other news drives the effect. The literature has traditionally focused on FOMC announcements to construct the shocks. But more recently, researchers (Swanson (2023), Swanson and Jayawickrema (2023)) have extended the approach to include other events, like speeches by Fed Leadership, FOMC press conferences, or Fed Meeting Minute releases. These papers find that other events can have at least as much effect on financial markets as FOMC announcements.

We follow Jarociński and Karadi (2020) and disentangle the pure monetary policy component of a monetary policy announcement from an "information" component. The former

can be thought of as a traditional monetary policy shock, where monetary policy itself is unexpectedly and exogenously tighter. The information component captures the situation when a policy announcement provides new information on (the central bank's view onf) fundamentals of the economy to the market. Jarociński and Karadi (2020) propose disentangling the two effects by their differential impacts on interest rates and equity prices, and present evidence that doing so improves the quality of the monetary policy shocks.

3.2 Data

Our approach requires high-frequency data on interest rates and the stock market. For the former we use the price of "fourth", or three-month ahead, Federal Funds Futures (FF4), as in Gertler and Karadi (2015); for the latter we use the S&P 500 index, as in Jarociński and Karadi (2020). For January 2008 to December 2023, we use tick-level data purchased from Bloomberg. We compute "raw" changes in these series over windows which start ten minutes before FOMC policy announcements and end twenty minutes after the announcement, as common in the literature. To extend our sample further back, we use data from Gürkaynak et al. (2022) which contains the same changes for these two securities, going back to the late 1980s. This data ends in June 2019, and we use the overlapping sample period to confirm that the raw changes agree. Appendix A.1 presents full details of how exactly we compute the relevant changes using the high-frequency financial data.

Following Swanson (2023) and Swanson and Jayawickrema (2023), we also develop a secondary shock series, which adds changes in market prices in windows around speeches by policymaker. However, data limitations mean that this secondary series is only available back to 2010. We discuss this series more in Appendix A.2.

3.3 A rotational decomposition

To decompose high-frequency movements in interest rates and equity prices into a pure monetary shock and a Central Bank information shock, we use a rotational decomposition. The advantage of this approach is that it relies only on the high-frequency financial market data and not on other macro time series. This is important for us since we want to use macro data to formulate a null distribution for the shocks in the test period using the filter. The test statistics for the hypothesis tests will use the true high-frequency shocks. If they also depend on the macro data then they are not independent of the relationship between the shocks and the data and so the tests will be invalid. In contrast, Jarociński and Karadi (2020) identify their shock as a sign restriction in a VAR which includes macro data, which would be problematic given the exercise we have in mind. In practice, though, our shock series and theirs are very similar.

We let the vector m_t be the changes in interest rates and the S&P 500 in the window around FOMC announcements in each month. We write

$$m_t = \begin{bmatrix} \Delta i_t \\ \Delta q_t \end{bmatrix}$$

We assume that the data has already been transformed to be mean zero and denote the variance of m_t by Γ.

We want to give the high-frequency data m_t a structural interpretation, converting it into structural shocks v_t, which are related by $m_t = Hv_t$ where v_t is mean zero and $\mathbb{E}v_t v_t' = I_2$. In particular, if the first element of v_t exhibits a negative correlation of interest rates and equity prices and the second a positive correlation, then we can interpret the elements of v_t as a pure monetary shock and a Fed information shock respectively. This imposes a sign restriction on H:

$$H = \begin{bmatrix} + & + \\ - & + \end{bmatrix}$$

We satisfy these restrictions with an orthogonal decomposition, imposing a parametric structure on the H matrix:

$$H = \begin{bmatrix} h_1 \cos\psi & h_2 \sin\psi \\ -h_1 \sin\psi & h_2 \cos\psi \end{bmatrix}, \qquad h_1, h_2 > 0 \quad \psi \in (0, \pi/2)$$

The intuition for this setup is that ψ is the angle of a line through $(\Delta i_t, \Delta q_t)$ below the x axis. This is the prime axis of rotation. It is close to, but not necessarily identical to, the OLS line of best fit. The two components of v_t then contain instructions for getting to the corresponding point m_t, expressed as a movement along the line of best fit (the first component) and perpendicular to it (the second component).

Solving for h_1, h_2, ψ such that $HH' = \Gamma$, we get that:

$$\begin{aligned} \tan 2\psi &= \frac{2\gamma_{12}}{\gamma_{22} - \gamma_{11}} \\ h_2^2 &= \frac{1}{2}\left(\gamma_{11} + \gamma_{22} + \frac{\gamma_{12}}{\sin\psi\cos\psi}\right) \\ h_1^2 &= \gamma_{11} + \gamma_{22} - (h_2)^2 \end{aligned}$$

where γ_{ij} is the (i,j)th entry of Γ.[9] We then create the structural monetary policy shocks from $v_t = H^{-1}m_t$.[10]

[9] This is similar to the decomposition in Jarociński (2022). However, there the decomposition does not impose unit magnitude on the v_t shocks and so cannot solve uniquely for ψ.

[10] Since $|H| = h_1 h_2 > 0$, H is always invertible

3.4 The shocks

We first take a look at the raw shocks, which are the changes in the fourth Federal Funds Future and the S&P 500 stock market index over the 30 minutes event window.[11] Summary statistics are displayed in Table 1. The total number of observation for FF4 is 298, and for S&P 500 is 297. The shocks are small: the mean absolute change is 2.77 basis points for FF4 (36.1 basis points for S&P 500) over the full sample, with a standard deviation of 5.48 for FF4 (55.8 for S&P 500). The mean absolute change is slight smaller since the pandemic, with 1.55 basis points for FF4, and 34.3 basis points for S&P 500. Our sample includes 31 FOMC meetings since March 2020, which we consider to be the point in time in which the COVID-19 pandemic started in the U.S.

Statistic	FF4			SP500		
	All	<03/2020	≥03/2020	All	<03/2020	≥03/2020
Mean	-1.21	-1.39	0.35	1.51	1.78	-0.77
Mean (Absolute)	2.77	2.91	1.55	36.06	36.26	34.34
Standard Deviation	5.48	5.71	2.31	55.76	56.72	47.46
Minimum	-37.00	-37.00	-6.00	-188.03	-188.03	-136.58
Maximum	12.00	12.00	5.50	407.58	407.58	71.32
Observations	298	267	31	297	266	31
Obs > Zero	94	81	13	149	131	18
Obs = Zero	77	67	10	3	3	0
Obs < Zero	127	119	8	145	132	13

Table 1: Summary Statistics for raw shocks for Federal Funds Future (FF4) and S&P 500 stock index (SP500), both in basis points.

The raw shocks are the change in the price of FF4 and the basis point change in the S&P 500 stock index over the 30 minute event windows. "All" refers to full sample ranging from January 1990 to January 2024, "< 03/2020" to before March 2020, and "≥ 03/2020" to since March 2020.

Figure 2 shows the cross-sectional plot of the raw shocks, m_t, over the estimation period. The solid line shows the first column of H. Each point can be uniquely decomposed into a distance along this line and a distance perpendicular to it. These distances are the entries in v_t and represent our pure monetary policy shock and Central Bank information shocks respectively.

Figure 3 shows the time series for these shocks, together with the two shock series from Jarociński and Karadi (2020). The latter shocks end in June 2019, while our shocks include the pandemic-period and end in January 2024.[12] The shocks from both approaches overlap very well. The correlations coefficient on the overlapping sample is 87 percent for the

[11] We transform the price p of the Federal Funds Future to an interest rate by setting $p = 100 - R$, where R is the arithmetic average of the daily effective federal funds rates during the contract month.

[12] Data limitations prevent us from extending the shock series in Jarociński and Karadi (2020) exactly.

monetary policy shock and 77 percent for the Central Bank information shock, respectively. Note that some discrepancy is not surprising due to the different methodology: Jarociński and Karadi (2020) include macroeconomic variables when constructing the shocks, while our approach – intentionally – doesn't. The shocks series from both approaches show a higher variance around "crisis" periods like around 2001 after the bursting of the dot-com bubble, as well as the Great Financial Crisis of 2008 and 2009. A period of some divergence between the two approaches is before 1994, where our monetary policy shocks tend to be smaller, while the Central Bank information shocks tend to be bigger. This is also the time period where the FOMC did not regularly issue a press release.

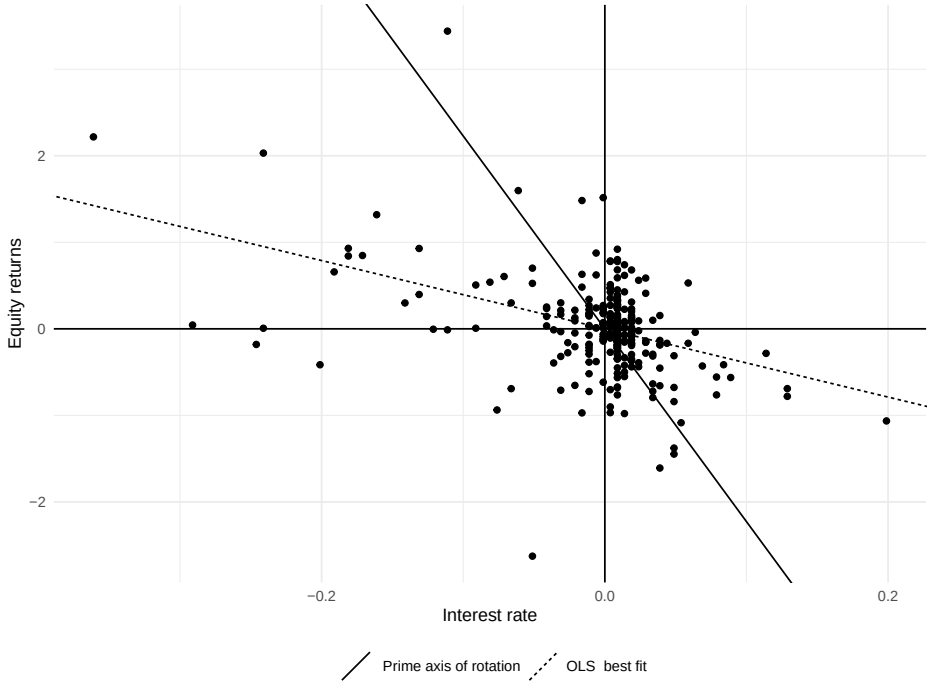


Figure 2: High frequency data: rotational decomposition

Figure shows changes in the interest rate and equity returns in thirty minute windows around FOMC meetings, aggregated monthly during 1992-2019. The prime axis of rotation corresponds to the vector in the first column of H, so ψ is the angle between it and the x axis.

4 Estimation and Validation

4.1 Data

We aim to incorporate information from a broad set of macro time series in our estimation, while still retaining interpretable impulse responses for comparison to those in the literature. To this end, we include both commonly-used macro time series in the vector y_t, as well as factors extracted from a wider set of variables. That is, we estimate a factor-augmented

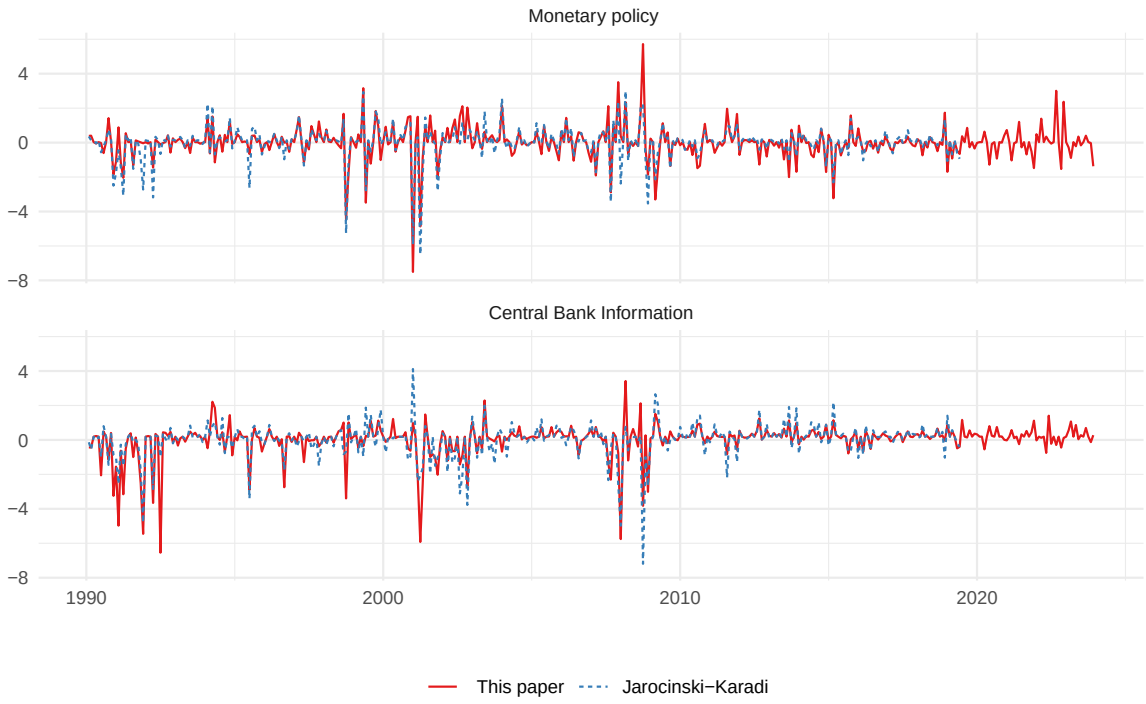


Figure 3: High-frequency shocks: comparison with Jarociński and Karadi (2020)

vector autoregression (FAVAR). In our baseline specification, y_t includes five variables directly – the Federal Funds Rate, log industrial production, the log of the Consumer Price Index, the employment-population ratio rate, and the Excess bond premium – as well as five factors summarizing further 18 variables covering financial markets, prices, real activity, and labor markets. These variables are listed in Table 2 and we use a principal components decomposition to create the factors. These five factors explain 70 percent of the variation in these covariates. Figure 4 plots the data for the five main series after detrending and deseasonalizing. The full data sample runs from January 1990 to December 2023, with data before December 2019 used for estimation and after January 2021 for out-of-sample testing. Figures A2 and A1 in Appendix A.3 plot the data including the factors and the extra data series respectively.

4.2 A Baseline View of the Transmission of Monetary Policy

We estimate the reduced form VAR in equation (1) on the 1990-2019 period, and interpret the first two rows of the Cholesky decomposition as the impact of shocks to monetary policy and central bank information. Although our focus is not on estimating these responses, the impulse responses are an important step in our analysis since they embody the view of the

Variable	Source	Transformation
Included directly		
Fed Funds Rate	Federal Reserve Board	Level
Industrial Production	Federal Reserve Board	Log Level
Consumer Price Index	Bureau of Labor Statistics	Log Level
Employment-Population Ratio	Bureau of Labor Statistics	Level, percent
Excess Bond Premium	Federal Reserve Board	Level
For inclusion as factors		
3 month treasury rate	Federal Reserve Board	Level
1 year treasury rate	Federal Reserve Board	Level
10 year treasury rate	Federal Reserve Board	Level
US dollar broad EER	Bank for International Settlements	Log Level
S&P 500 Index	Standard & Poors/Dow Jones	Log Level
CPI inflation	Bureau of Labor Statistics	Percent
PCE deflator	Bureau of Economic Analysis	Log Level
Real oil price, CPI deflated	Energy Information Administration	Log Level
Commodity Price Index	Commodity Research Bureau	Log Level
Federal Government Outlays	U.S. Treasury	Log Level
Federal Government Receipts	U.S. Treasury	Log Level
Unemployment Rate	Bureau of Labor Statistics	Level, percent
Average Weekly Hours	Bureau of Labor Statistics	Log Level
Average Hourly Earnings	Bureau of Labor Statistics	Log Level
Residential Investment	Bureau of Economic Analysis	Log Level
Manufacturing Orders	U.S. Census Bureau	Log Level
Consumer Confidence	Conference Board	Level
Real Personal Consumption	Bureau of Economic Analysis	Log Level

Table 2: Data and sources

All variables are monthly and deseasonalized, and detrended using a Hodrick-Prescott filter. The division of series into those included directly and those which are included as factors follows the baseline specification.

transmission of monetary policy to which we want to compare more recent data. We allow for $J = 2$ lags in our baseline specification, selected by the Akaike information criterion. A short lag length is not entirely surprising since the inclusion of the factors means we have many variables in the VAR – extra lags are more heavily penalized and there is less marginal information from extra lags. Longer lag lengths produce similar results (See Appendix B.2).

Figure 5 shows the response of the five headline variables to the monetary policy shock, normalized to a one percentage point increase in the Federal Funds Rate. Factors are included, but not shown; see Figures A3 and A4 in Appendix B.1 for the responses of the factors). The results are similar to those in the literature. Industrial production falls sharply in the months after the monetary shock, with a maximum contraction of around 3 percentage points. The real impact reaches its maximum around eight months after the shock although the effect persists for rather longer. We also estimate a decline in prices on impact, although this is a little less persistent than estimates elsewhere. The decline in real activity and

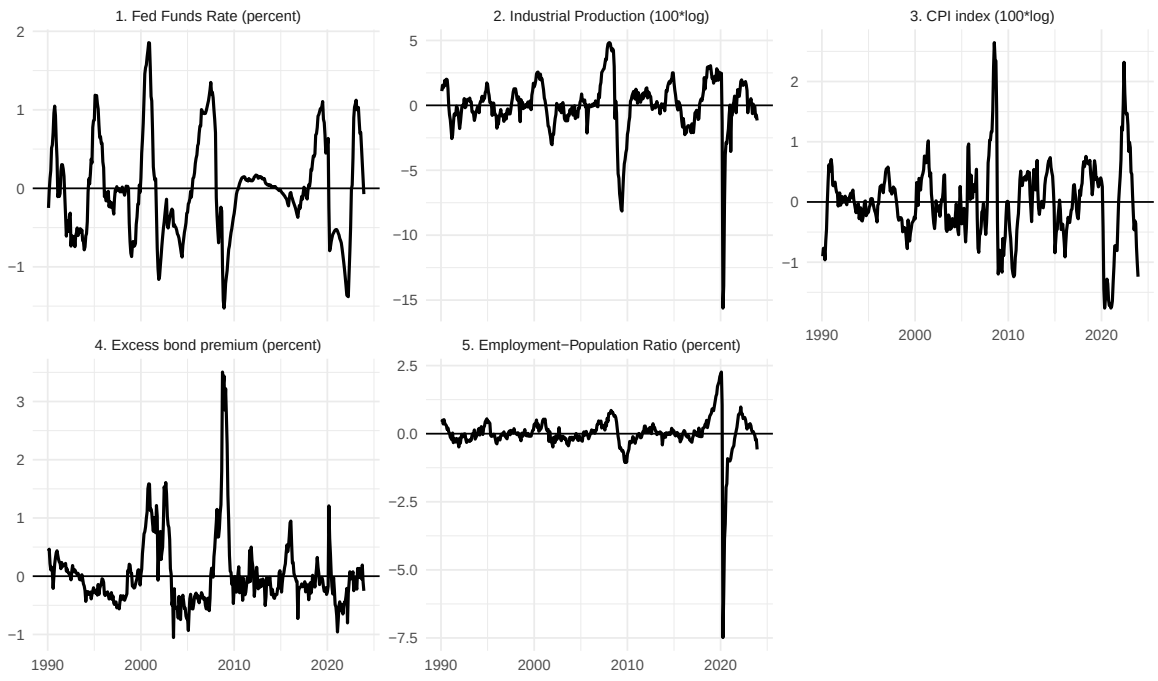


Figure 4: Time series data used in the baseline VAR

Figure shows the data series which are in the headline vector autogressions after transformations, excluding factors.

prices is also associated with a subsequent reduction in interest rates, consistent with an endogenous response to lower output and inflation.

Our estimated Central Bank information shock is also broadly consistent with other estimates, both qualitatively and quantitatively. Positive information about the future state of the economy leads to an increase in real activity and prices with interest rates rising in response and employment increasing.

In Figures 5 and 6 we scale the shocks such that the initial impulse for the Federal Funds rate is one percentage point. Typically, this requires a relatively large high-frequency shock; around 120 standard deviations of the monetary policy shock and 35 of the Central Bank information shock. This occurs because the identified shock captures only a small part of the variation of interest rates. As mentioned above, an alternative approach is to treat v_t as imperfectly correlated instruments for a broader set of monetary shocks. As Paul (2020) shows, this produces identical impulse responses, just rescaled – effectively a change-of-units to account for the fact that not all the monetary policy shocks are captured by the high-frequency series. We verify this in Figures A29 and A30 in Appendix B.5 where we present the point estimate for the impulse responses using the shocks as external instruments. As they should, they match the impulse responses in Figure 5 and 6 identically up to a scalar.

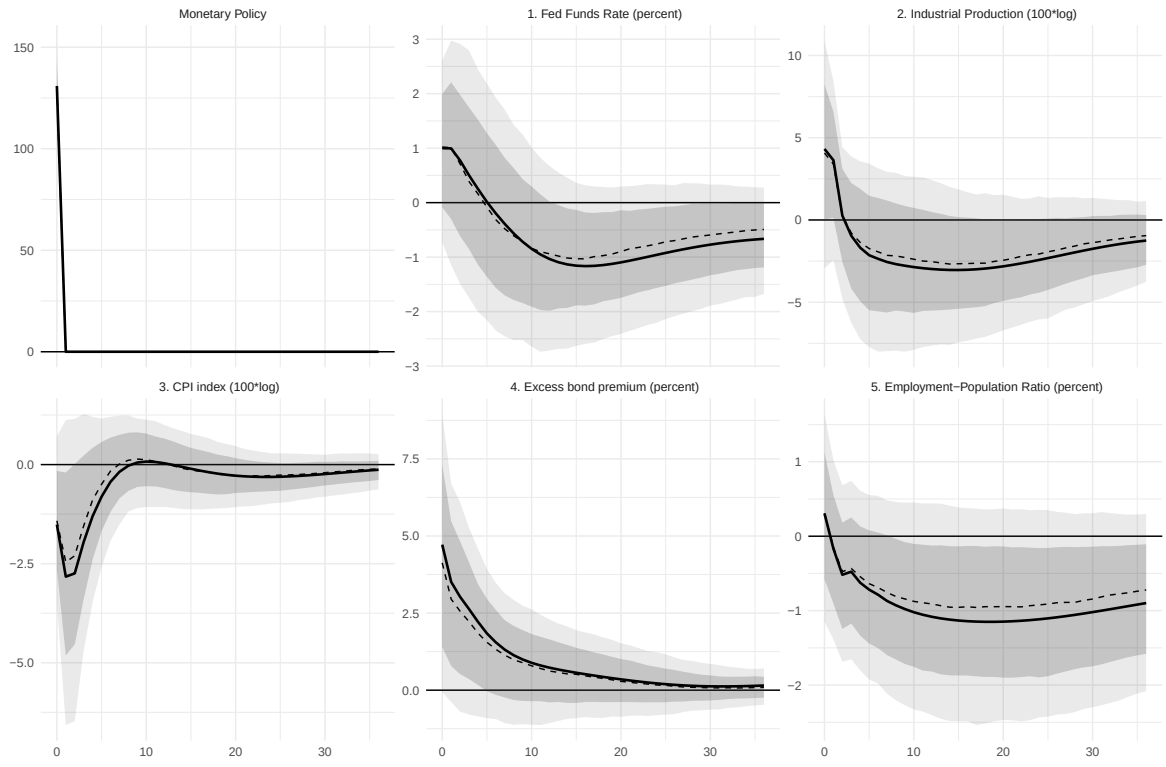


Figure 5: Impulse responses: Monetary policy shock

Figure shows the response of the headline variables to a monetary policy shock, scaled to a 1 percentage point increase in the Federal Funds Rate. The estimated VAR includes the five main variables, plus both shocks and five factors. Solid lines are point estimates and shaded regions are the 68 and 90 percent confidence intervals from a bootstrap with $K = 1000$ replications. Dashed lines show median responses from the bootstrap. Data is monthly and the estimation sample is January 1990-December 2019

4.3 Robustness

In Appendix B.2 we investigate the sensitivity of the estimated structural VAR model to a variety of alternatives. We show there that longer lag lengths make little difference to the impulse responses (Figures A5 and A6). Changing the interest series has little impact either – the responses using a one-year government bond rate in place of the federal funds rate are identical, and when using a three-month rate, only the interest rate response changes slightly (Figures A7 and A8). A similar lack of difference holds when variously replacing our headline series with other measures of inflation, activity, financial market variables, labor market activity, and residential investment (see Figures A9-A14). We suspect that this stability across alternative specifications is likely a result of the factors that we use to augment our VAR. Since they soak up much of the correlated variation in other series, replacing one or the other of the headline measures makes little difference to the estimated impulse responses.

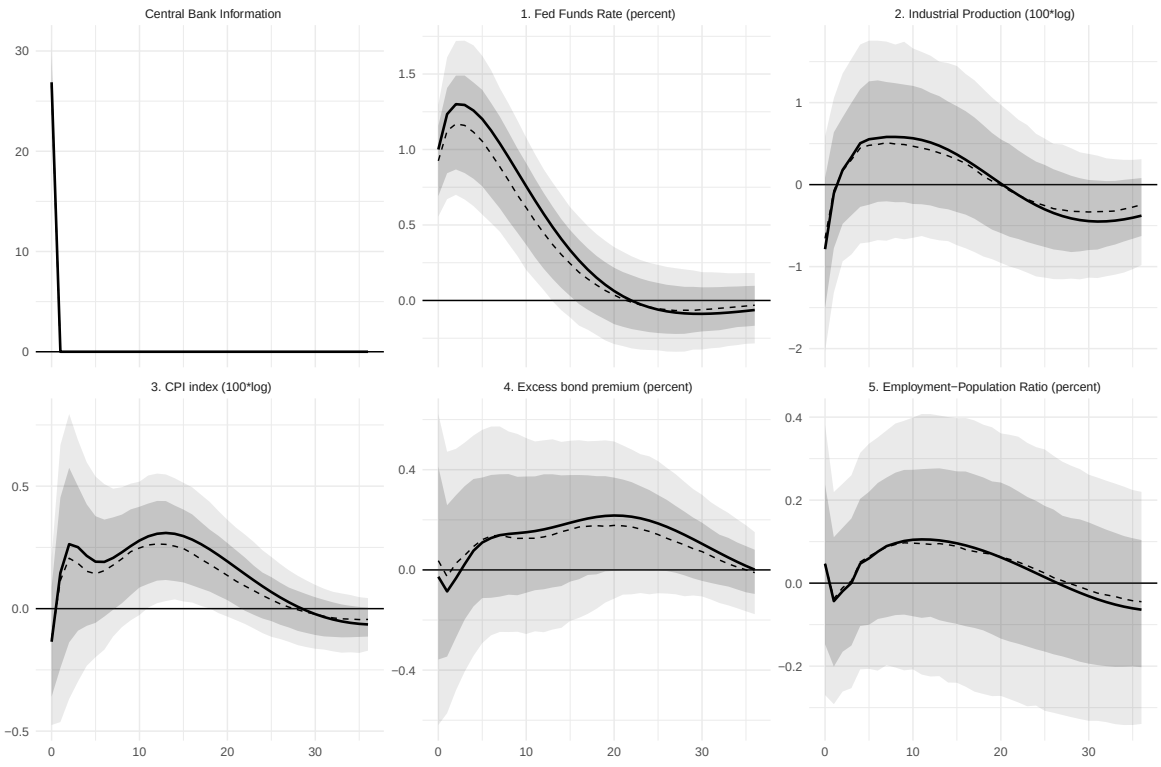


Figure 6: Impulse responses: Central Bank Information shock

Figure shows the response of the headline variables to a monetary policy shock, scaled to a 1 percentage point increase in the Federal Funds Rate. The estimated VAR includes the five main variables, plus both shocks and five factors. Solid lines are point estimates and shaded regions are the 68 and 90 percent confidence intervals from a bootstrap with $K = 1000$ replications. Dashed lines show median responses from the bootstrap. Data is monthly and the estimation sample is January 1990-December 2019

We verify the importance of the factors in Appendix Figures A15 and A16 where we show that removing the factors has a relatively large impact, with interest rates staying persistently higher and output and employment slower to respond. The inclusion of a relatively rich set of factors also largely explains why our impulse responses differ somewhat from canonical estimates such as those of Jarociński and Karadi (2020). There, the peak decline in industrial policy per unit increase in peak short rates for a monetary policy shock is about three, and is reached at a horizon of around 10-20 months. In contrast, our baseline results have a similar peak response of industrial production, but reach it much faster, at around 7 to 8 months. They also produce a hump-shaped response for short-term interest rates, which we do not. When we exclude the factors from our VAR, though, these differences largely go away. This is consistent with the notion that we are capturing similar shocks to Jarociński and Karadi (2020) (as Figure 3 suggests) but the inclusion of a wide set of factors means that we control for other sources of correlated variation. Figure A17 compares our results directly to the corresponding monthly impulse responses from the online Appendix of Jarociński and

Karadi (2020), showing that when we remove factors, match variables, and estimate on the same sample we recover very similar impulses for the monetary shock.

We also check that our results are robust to the extension of our shock series to include speeches since 2010, as suggested by Swanson (2023). Appendix Figures A24 and A25 present the corresponding impulse responses using this shock series and find that they are broadly similar, albeit with a slightly more delayed response of the real activity variables, industrial production and employment.

To be clear, the broader point of this section and the previous one is not to claim that we have the definitive measure of how monetary policy affects the macroeconomy. Rather, it is to convince the reader that the view of the monetary transmission mechanism against which we test the 2022 data is broadly sensible. And although reasonable people may differ in the details of exactly how policy transmits, the baseline view that our impulse responses represent is consistent with mainstream views.

4.4 Validation

To check that our method works, we run the filter on the estimation period (1990-2019), backing out the two implied policy shocks from the data alone. This, of course, should work—that is the point of a check—but it is also a chance for our method to fail in case there were something wrong with it.

Table 3 assesses the point estimates from our filter, regressing the actual shocks on our inferred shocks. If our filter is truly producing a conditionally unbiased estimate of the shocks then this regression should have a slope of one – a unit change in the estimated shock predicts a unit change in the actual shock – and have an intercept of zero – no bias. This holds true for both shocks. One less obvious point is that the R^2 is low for both shocks, especially the pure monetary policy shock. This says that much of the variation in the realized shocks cannot be captured by our VAR. One possibility is that our model is mis-specified, but the other is simply that this might be unavoidably difficult. In particular, if other shocks are driving much of the variation in macro outcomes it will be hard to infer the monetary shocks from macro data.

Whatever the degree of difficulty in recovering shocks from macro variables, our measure of uncertainty should always reflect this. To check this, Table 4 reports the coverage ratios of the realized shocks computing percentile ranges at each point in the estimation sample using the mixture of normals in equation (8). If our assumptions on the distribution of δ_t and subsequent calculations are correct, then the true shocks should be inside the $x\%$ confidence interval x percent of the time. The column labelled "Data" computes the coverage for the actual data, and suggests that the filter characterizes the 90 and 95 percent intervals fairly well, but perhaps a little wide for the middle of the distribution. This is consistent with the

	Monetary Policy	Central Bank Information
	(1)	(2)
Estimated shock	1.011***	1.005***
	(0.187)	(0.122)
Constant	0.005	−0.025
	(0.051)	(0.049)
Observations	356	356
R^2	0.076	0.161
Adjusted R^2	0.074	0.158
Residual Std. Error (df = 354)	0.966	0.921
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 3: Model validation: Regressing actual on estimated shocks

Regression of actual shocks on the shocks estimated via the filter, monthly data 1992-2019. Spherical standard errors shown in parentheses. This is a test of our method, and if it works should return coefficients of one and zero on the estimated shock and the constant respectively.

identified shocks being highly leptokutic, with excess kurtoses of around 13.

To understand how this result might be affected by the particular realization of the data, we supplement this calculation with two simulation exercises. We first use the estimated data generating process to simulate 500 samples of the same length as the true data, and compute coverage ratios for each. These make precise the idea that the filter does a good job of measuring the tails of the distribution of shocks, but in the middle of the distribution produces confidence intervals which are a little too wide. Specifically, the 95 percent coverage ratios seem well-centered (medians are close to 95 percent) and informative (90 percent of simulated ratios are within 2 percentage points of the truth) for both shocks. For the 90 percent confidence intervals performance is only fractionally worse and still quite good for the monetary policy shock. But the 68 percent confidence interval seems to be systematically too wide. This means that any tests using this part of the distribution will under-reject data generating process different to the mull. We return to this point later.

We also generate a single very long simulation of 12,000 observations (i.e. 1,000 years worth of data) and compute the corresponding in-sample coverage ratios. This should give a sense of whether the confidence intervals are consistent, and the extent to which any biases might ease in large samples. In the end, the corresponding coverage ratios look very similar to the data, implying that we have sufficient data to rely on our conclusions.

Shock	Interval	Data	Short simulation percentiles					Long Simulation
			2.5	5	50	95	97.5	
Monetary Policy	68 Percent	80.5	76.1	76.6	80.3	83.9	84.5	80.5
	90 Percent	91.6	88.8	89.3	92.7	94.9	95.5	91.6
	95 Percent	94.5	92.1	92.7	95.2	97.1	97.2	94.5
Central Bank Information	68 Percent	84.7	79.7	80.3	84.2	87.8	89.0	84.7
	90 Percent	93.7	90.7	91.3	93.2	95.8	96.1	93.7
	95 Percent	94.8	92.7	92.7	94.9	96.8	96.9	94.8

Table 4: Model validation: Coverage ratios

Column labelled "Data" reports the in-sample coverage ratios using a VAR estimated on the 360 periods in the estimation period. The columns labelled "Short simulation percentiles" report the percentiles of the same coverage ratios but using 500 simulated samples with the same length as the data. Finally, the column labelled "Long Simulation" reports the results using a single simulation of 10,000 periods. All calculations define the coverage ratio as the fraction of observations falling inside the central 68, 90, or 95 percent confidence interval for the filter. Throughout, confidence intervals are computed using a bootstrap with 1000 draws.

5 Results from the Filter

5.1 True versus filter-implied shocks

We apply the VAR-based filter to the post-2021 data. This calculates in each period the distribution of monetary (and other) shocks under the null hypothesis that the data generating process has remained unchanged. By then comparing the realized monetary shocks from the high frequency data to this sequence of distributions we perform an out-of-sample test that the data generating (and thus the monetary transmission process) remains unchanged. If the data generating process is unchanged, the realized shocks should be statistically indistinguishable from the inferred ones.

Figure 7 presents the headline results from this experiment. The shaded regions show the confidence intervals for shocks conditional on the data under the null that the data generating process is unchanged, computed using the distribution in equation (8). The blue line shows the shocks measured from the high frequency financial market data. Both series show notable and persistent departures of the realized shocks from the inferred ones. In particular, the measured monetary policy shock exceeds the point estimate from the filter throughout January to July 2022. In other words, the FOMC was repeatedly moving financial markets in ways consistent with tightening but the macro data were behaving as if they were loosening. That is, the transmission of monetary policy was looser. Likewise, the Fed information shock exceeds the point estimate From April 2021 until April 2022. This means that markets were interpreting Fed announcements at the time as if they were revealing positive information about the state of the economy, but the monthly macro data did not subsequently confirm that.

How sure can we be of these conclusions? The shaded regions in Figure 7 provides a metric to answer this question period-by-period, as they represent the distribution for the data-consistent shocks, under the null that the transmission of monetary policy is unchanged. As should be apparent from the figure, the evidence in any given month for either shock transmitting differently from usual is not strong – with a few exceptions, the blue dots are typically inside the center of the distribution. Bear in mind, though, the evidence of the previous section that the central confidence intervals are too wide. And so true statistical power is likely understated here. And so the statistical confidence that monetary policy transmission is weaker in any given period is typically low. However, the measured monetary policy shock is not in the upper part of the filter distribution in just one period, but rather for sequence of periods. Even if we cannot be very confident that the transmission of monetary policy is weaker in any *given* period, perhaps we can pool evidence across these periods to conclude that monetary policy transmission was weaker on average, or at least some of the time. Much of the rest of this paper attempts to tackle this issue.

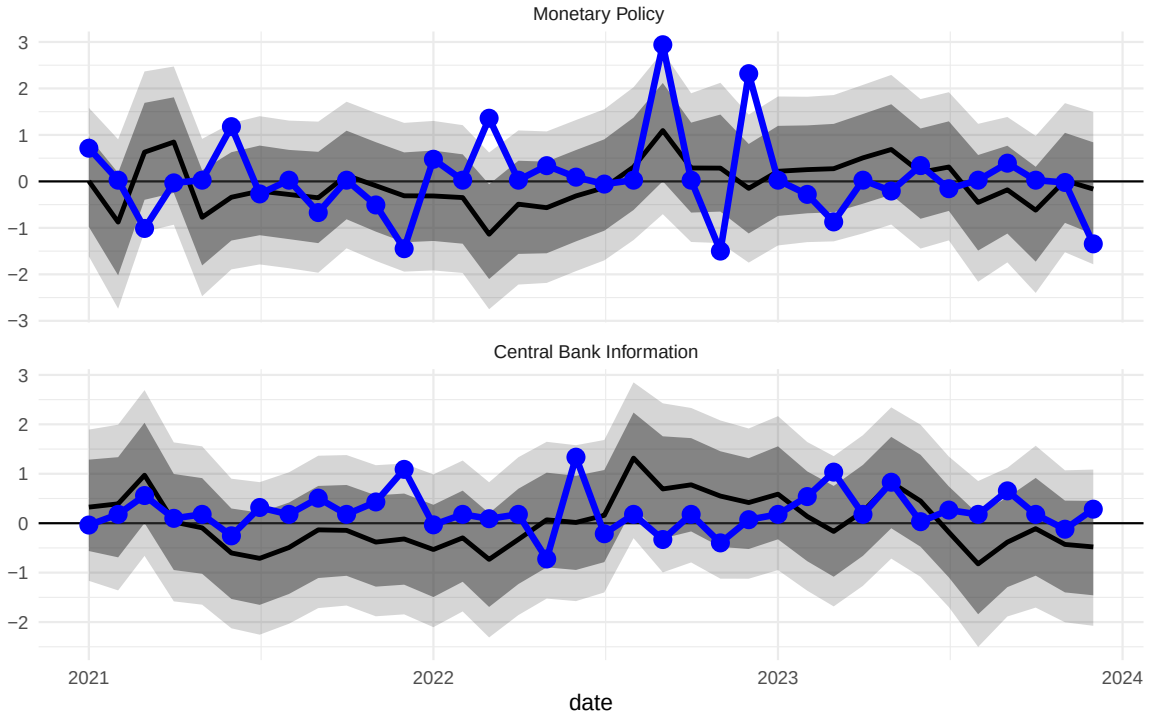


Figure 7: Inferred and realized shocks: 2021-2023

Solid lines are point estimates of the inferred shocks under the null hypothesis that the data generating process remains unchanged. The shaded regions are the 68 and 90 percent confidence intervals from a bootstrap with $K = 1000$ replications. The blue line with large dots is the actual shock, computed from high frequency data.

Shock	Month	Consecutive periods above	Consecutive periods below
Monetary Policy	2007-03-01	0	6
Monetary Policy	2008-04-01	7	0
Monetary Policy	2020-06-01	0	9
Monetary Policy	2022-07-01	7	0
CB Information	2010-09-01	9	0
CB Information	2014-12-01	9	0
CB Information	2019-04-01	10	0
CB Information	2022-04-01	13	0

Table 5: Historic consecutive periods of unusual shocks: Full sample, 1990-2023

Table shows terminal months of the longest streaks of high-frequency observed shocks above or below the point estimate. Episodes included in the table are the four longest consecutive sequences above or below the point estimate for each shock. For example, all other observations have 6 or fewer consecutive observations above or below the inferred monetary policy shock.

5.2 Informal assessment

We first start with an informal test of this idea, asking: how rare is it that the observed shocks are persistently above or below the filter-implied estimates? It is very rare. For both shocks, the periods are the (sometimes joint) longest streaks where the observed data exceed the point estimate from the filter in the entire sample, starting in 1990. Table 5 puts these streaks in context, reporting the longest streaks where the observed high-frequency shocks are consistently above and below the filter-implied point estimates, across the whole sample (1990-2023). This shows that only in times of extreme stress—the Global Financial Crisis and the onset of the COVID-19 pandemic—have there been even a comparable number of consecutive observations of the monetary policy shock above or below the filter point estimates.

5.3 Formal assessment: Hypothesis testing

We can also calculate a more rigorous assessment of the extent to which these shocks are unusual. Given that we have explicit distributions for the shock conditional on the data and the data generating process in each period, we can compute p-values for the joint hypothesis test that a whole sequence of shocks is drawn from the filter-implied distribution.

Formally, we denote by η_t the difference between the observed high-frequency shocks and the mean of the filter-implied estimate, so $\eta_t = v_t - \bar{\delta}_t$. We then define $\mu_t = \mathbb{E}\eta_t$ as the mean of η_t, which is a parameter. If the relationship between the high-frequency shocks and the macro data is unchanged then this will be zero in every period. Based on the preceding section, we have reason to believe that this may not be true during parts of 2021 and 2022. And thus for any particular period $\mathcal{T}$ we test three possible hypotheses about each shock i:

Test 1 $H_0 : \mu_t^i \leq 0$ for some $t \in \mathcal{T}$ vs. $H_1 : \mu_t^i > 0$ for all $t \in \mathcal{T}$

Test 2 $H_0 : \mu_t^i \leq 0$ for all $t \in \mathcal{T}$ vs. $H_1 : \mu_t^i > 0$ for some $t \in \mathcal{T}$

Test 3 $H_0 : \bar{\mu}^i \leq 0$ vs. $H_1 : \bar{\mu}^i > 0$, where $\bar{\mu}^i = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \mu_t^i$

Intuitively these three tests capture the idea that in a given period, the transmission of monetary policy is different 1) at all points during $\mathcal{T}$, 2) at least some of the time during $\mathcal{T}$, and 3) on average across $\mathcal{T}$. All three tests are explicitly one-sided, which is important in our application since we have a priori evidence that μ_t has the same sign for multiple consecutive periods. However, this makes the rejection regions and corresponding p-values a little harder to compute than in the two-sided equivalents. Note also that in only test 3, where we compute a sample mean, is a central limit theorem likely to apply directly. In tests 1 and 2 we add an extra parameter with each extra observation, and so the power of the test does not grow with the number of periods in $\mathcal{T}$. Indeed, test 1 is a very stringent test, and the null will be rejected only if we have strong evidence that transmission is changed in all periods.

Shock	$\mathcal{T}$ Start	$\mathcal{T}$ End	$\lvert\mathcal{T}\rvert$	Test 1	Test 2	Test 3
Monetary Policy	2022-02-01	2022-02-01	1	0.349	0.346	0.349
	2022-02-01	2022-03-01	2	0.349	0.012**	0.017**
	2022-02-01	2022-04-01	3	0.349	0.021**	0.021**
	2022-02-01	2022-05-01	4	0.349	0.025**	0.013**
	2022-02-01	2022-06-01	5	0.349	0.036**	0.014**
	2022-02-01	2022-07-01	6	0.468	0.054*	0.021**
Central Bank Information	2021-05-01	2021-05-01	1	0.383	0.382	0.383
	2021-05-01	2021-06-01	2	0.383	0.537	0.316
	2021-05-01	2021-07-01	3	0.383	0.346	0.149
	2021-05-01	2021-08-01	4	0.383	0.363	0.103
	2021-05-01	2021-09-01	5	0.383	0.379	0.074*
	2021-05-01	2021-10-01	6	0.383	0.442	0.072*
	2021-05-01	2021-11-01	7	0.383	0.407	0.046**
	2021-05-01	2021-12-01	8	0.383	0.233	0.017**
	2021-05-01	2022-01-01	9	0.383	0.260	0.014**
	2021-05-01	2022-02-01	10	0.383	0.289	0.013**
	2021-05-01	2022-03-01	11	0.383	0.270	0.008***
	2021-05-01	2022-04-01	12	0.383	0.300	0.007***

Table 6: p-values for joint hypothesis tests

Table shows p-values for the joint hypothesis that consecutive observations are drawn from the null. Each line considers three tests that the observed high frequency shocks are drawn from a distribution with a higher mean than would be expected based on pre-COVID transmission of monetary policy, over a period of consecutive months. $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

The top half of Table 6 reports the p-values for these tests, evaluated by likelihood ratio,

for a series of windows starting in February 2022 for the monetary policy shock.[13] The values in each row correspond to the maximum probability of falsely rejecting the H_0 under the null for a sequence of policy shocks starting February 2022. As expected, the p-values for test 1 suggest that we cannot conclude that the relationship between the high frequency shocks and the data is always different in any part of the first half of 2022. This is because there is always at least one period where the observed high frequency shocks are not sufficiently different from the filter-implied estimates to convincingly reject the null. However, tests 2 and 3 provide stronger evidence that monetary policy is weaker than would normally be expected. They show that there is evidence at around the 1 to 3 percent confidence level that monetary policy is weaker at least some of the time, as well as on average.[14]

The lower half of the table repeats a similar exercise for the central bank information shock. There, there is less precision about whether the relationship between this shock and the macro data is different in any one period. However, in the longer test periods, the power of the central limit theorem starts to show. If we reject the null of unchanged average transmission of the central bank information shock during the second half of 2021, we would do so incorrectly less than 1 percent of the time. In other words, we can be very confident that the impact of the central bank information shock was on average different during this time.

5.4 The role of uncertainty

The measured uncertainty in our estimates is essential to computing the test statistics for the hypothesis tests above. So what drives them? In Figure 8 we decompose the 90 percent confidence intervals from Figure 7 into their contributions from model uncertainty (i.e. that coming from the fact that the true model parameters are unknown) and from the inference uncertainty (i.e. that coming from the filter). The latter is constant, since it reflects the lack of knowledge over the shocks due to missing information, which does not change. Most of the time, this is the main source of uncertainty, and arises because the shocks are fundamentally hard to disentangle. In some periods, however, the contribution from the model uncertainty is substantial. This is notably true in early 2021, and occurs when the reduced form shocks are largest, since uncertainty over the parameters is amplified by the magnitude of the shocks. This also explains why the contribution from model uncertainty is low in other periods, when smaller shocks imply less of an impact from model uncertainty.

[13] Appendix C describes in detail how these are computed. Note that the excess width of the central confidence intervals shown in Table 4 means that the computed p-values are likely upper bounds on the true values.

[14] It might seem odd that in some cases the p-value for test 3 is less than that for test 2. After all, how can one be more confident that the average of a set of values is positive than that any individual observation is? The reason is that the rejection region for test 2 is typically larger than that for test 3, and so the p-value is higher. More intuitively, rejecting test 2 requires only providing evidence that a single v_t^i is sufficiently large for some $t \in \mathcal{T}$. With large enough $\mathcal{T}$, the likelihood of at least one freakishly large draw becomes quite large *even under the null*, and so the p-value for any given observation can be quite high.



Figure 8: Uncertainty decomposition: 2021-2023

Figure shows decomposition of 90 percent confidence intervals into their contributins from model uncertainty, inference uncertainty, and their interaction.

5.5 Robustness

In Appendix B.3 we repeat these hypothesis tests for the monetary policy shock using the same set of alternative specifications as for the impulse responses. In Figures A18-A23 we find that the results are minimally affected by changes in lag length, changing the short-term interest rate series, replacing the excess bond premium with the exchange rate or long-term interest rates. Changing the number of factors has some effect, generally strengthening our key result that the transmission of policy was lessened on average. Likewise, replacing the inflation measure with core PCE or including residential investment tend to (weakly) strengthen this result. The only alternative specification which weakens this result at all is where we change the labor variable, to either the unemployment rate or aggregate hours.

In Appendix B.4 we repeat the hypothesis tests using the shock series which includes policymakers' speeches and show that they strengthen our results. In Appendix Figure A26 we show the post-2020 outturns compared to the filter-inferred shocks. The differences are larger than in the baseline, a fact confirmed by Appendix Table A1, which rejects unchanged monetary transmission even at the 1 percent level. Similarly, Figures A27 and A28 imply that monetary policy is statistically significantly weaker even if the covariance of other shocks is different.

6 Interpretation

In the preceding section, we evaluated the strength of the evidence that the link between the data and high-frequency identified shocks had changed sometime in 2022. But what does it really mean to reject any of the tests laid out above? One obvious interpretation is that it means that the transmission of monetary policy is somehow different. But it does not tell us how much it differs, or in what ways, nor if other interpretations might be equally valid. This section thus attempts to answer three related questions. Quantitatively, how much weaker would the transmission of monetary policy have to be to reconcile the observed high-frequency shocks with the data? What other factors, such as the changing impact of other shocks, could explain the results of the previous section? And what economic forces might have caused any such weakness?

6.1 Set up

The key ingredients in our filter which maps data into implied structural shocks are the matrices $\hat{D}_y$ and $\hat{\Omega}_y$, which together define the mean and covariance of the filter estimate conditional on the parameters, $\delta_{t|t,\hat{\theta}}$.

We can use modified $\hat{D}_y$ matrices to consider what would happen if the null transmission of the shocks were different by multiplying by a diagonal matrix of positive numbers, Λ, which we partition as:

$$\Lambda = \begin{bmatrix} \Lambda_m & 0_{K\times N} \\ 0_{N\times K} & \Lambda_m \end{bmatrix}$$

Then the modified causal impact matrix is:

$$\tilde{D}_y = D_y\Lambda = [D_m\Lambda_m D_x\Lambda_m]$$

This represents the statistical model where the impact of each of the shocks has been scaled by the corresponding entry of Λ.

For example, if top-left first entry of Λ were 0.9 and all other entries were 1, this would capture the scenario where the transmission of the pure monetary shock was uniformly ten percent less effective, in a specific and well-defined sense: a given monetary shock moves all outcomes by ten percent less than in the baseline model. Running the filter on the same data but replacing D_y with $\tilde{D}_y$ would then produce a distribution for the high frequency indicators under the new null, one where the impact of monetary shocks was ten percent weaker than usual. Performing our regular hypothesis tests using this modified model would then permit a test of the hypothesis that monetary policy was at least ten percent weaker on any given sample.

A natural question to ask is how one should pick the entries of such a modifying Λ? After all, if we alter some column of the impact matrix D, then without an offsetting modification of the other shocks the data generating process will no long have the same covariance as the data. In the example above, proposing that the monetary shocks have ten percent less impact than usual will mean that some of the variation in the data otherwise explained by the monetary shock is now unexplained by any shock in the model. However, because the other shocks matter only to the extent that they affect the covariance of the data, $\hat{\Omega}_y$,[15] one need only be explicit about the assumptions on these other shocks *up to the extent that they affect this covariance*. This allows for something of an end-run around the problem of missing variation. For example, if one posits a modified shock matrix $\tilde{D}_y$ where only the monetary shock is affected, then computing the filter with $\hat{\Omega}_y$ unchanged will recover the distribution of monetary shocks under the assumption that the unidentified shocks change *in whatever ways are needed* to guarantee that the data-generating process still matches that estimated from the sample. The cost of this shortcut is that the time series distributions for the other, unidentified, shocks will not be correct. But we are not interested in those. Of course, if one thinks that the variance of the data has changed on the test period, one can replace $\hat{\Omega}_y$ with something different. We return to this point later, in Section 6.3.

6.2 The changing impact of monetary policy

We start by trying to quantify the change in the strength of the transmission of monetary policy, assuming that the residual covariance in the post-2022 period, $\hat{\Omega}_y$ is unchanged. To do this, we consider a series of modified structural impact matrices:

$$\tilde{D}(\lambda_1) = D_y \times \begin{bmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & 1 & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{bmatrix}$$

for $\lambda_1 \in (0, 1)$. Then $\tilde{D}(\lambda_1)$ embodies the idea that monetary transmission is only λ_1 as strong as in the pre-COVID baseline. Put differently, interest rates would need to increase by around $1/\lambda_1$ percentage points to have the same effect as a 1pp increase pre-COVID.

For each λ_1 we then re-run our hypothesis tests from the preceding section, replacing D_y with $D_y(\lambda_1)$. Given the excessive stringency of test 1, which never rejects the null, we drop it in this and future exercises. Figure 9 plots the resulting p-values for the six monthly windows $\mathcal{T}$ starting in February 2022. The connection to Table 6 is that the rightmost values (i.e. at $\lambda_1 = 1$) match the numbers in the corresponding columns of Table 6.

[15] See equation (7) and subsequent discussion.

Overall, this plot suggests that the efficacy of monetary policy dropped quite considerably in March and April 2022 – during this time one would struggle to reject at the 5 percent level any null where monetary policy was stronger than about 75 percent of its usual efficacy. In other words, for every three interest rate hikes that the Fed did during this period, they would have needed to do approximately one more to deliver the same degree of tightening as one would have expected pre-COVID. Later in the year, this statement continues to hold true on average. That said, this figure also makes clear that the evidence we present has some limitation. At a 1 percent significance level, we would conclude that monetary policy was incrementally weaker in just April.

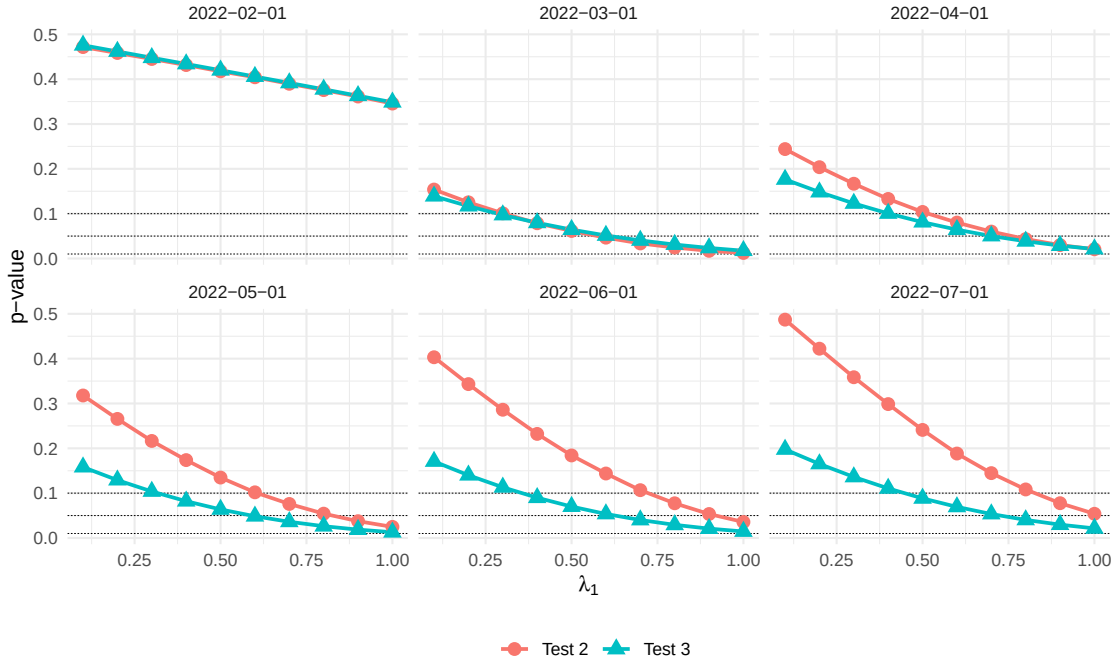


Figure 9: p-values for hypothesis tests as strength of monetary policy shock varies

Figure shows p-values for hypothesis tests described in Section 5.3, under a series of nulls, indexed by λ_1. Horizontal lines show 1, 5, and 10 percent confidence levels. The interpretation of λ_1 is that it represents a null where monetary policy is only λ_1 as effective as pre-COVID. Each panel computes the corresponding hypothesis tests on a window starting in February 2022 and ending in the titular month. The values for both lines when $\lambda_1 = 1$ thus match the corresponding columns in Table 6.

Because other shocks only matter to the extent that they affect the covariance of the reduced form residuals, these results hold identically no matter how other non-monetary shocks might have changed so long as the *covariance* of the data is unaffected. In the next section, we relax that assumption.

6.3 Allowing for changing variances

It is, of course, possible that the nature of other shocks changed during 2022. After all, the post-COVID period was a highly unusual one with many of the after-effects of a once-in-a-lifetime global pandemic still being worked out. One of the advantages of our approach is that we can capture the impact of such changes on our filter-implied monetary shocks without specifying exactly how these other shocks changed. This is fortunate, since we do not identify them. Instead, it is enough to describe solely how any such changes would have affected the reduced form residual covariance matrix, $\hat{\Omega}_y$. Indeed, if the reduced form residual covariance matrix had changed during the test period, this would undermine that interpretation of our results as capturing a change in the efficacy of monetary policy.[16]

For example, because the filter expects that tighter monetary policy will, at the margin, produce lower inflation – see Figure 5 – it will interpret the persistence of inflation in spite of high interest rates as a sign of weaker monetary policy. Of course, other offsetting shocks could be driving the slow decline in inflation, but the filter adjusts for this based on the usual correlation of the other macro variables. However, if the variance of those offsetting shocks has changed then the filter will account for them incorrectly.

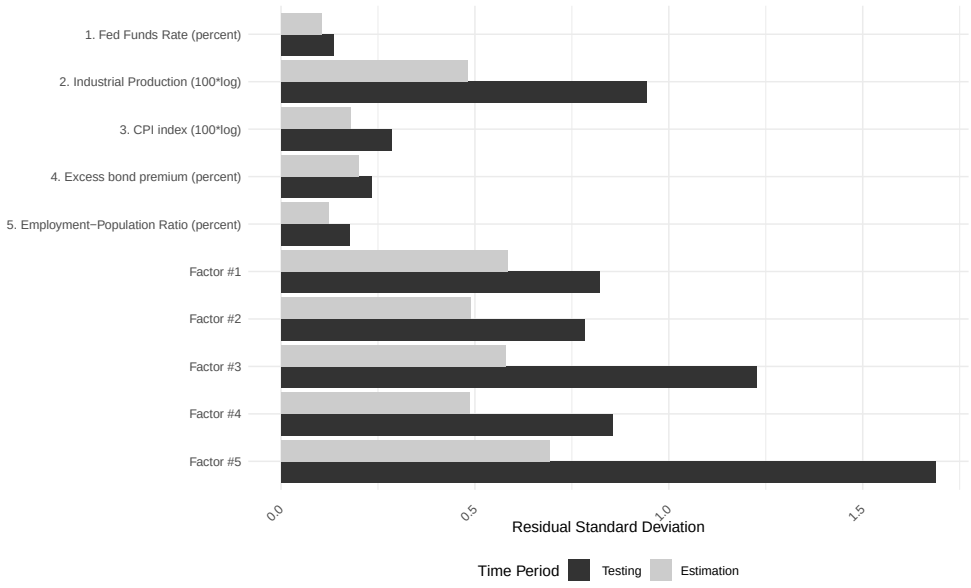


Figure 10: Reduced form residual standard deviations

Bars show the standard deviation of the reduced form shocks on the estimation and testing periods. In both cases, the autoregressive parameters, B_j, are assumed to be those estimated on the pre-COVID data.

Figure 10 presents preliminary evidence that the reduced form covariance has indeed changed. It shows the standard deviations of the reduced form residuals for the data series

[16] A further possibility that we do not investigate is that the autoregressive coefficients of the VAR also changed. The brevity of the post-tightening period means that such an exercise will likely be highly imprecise.

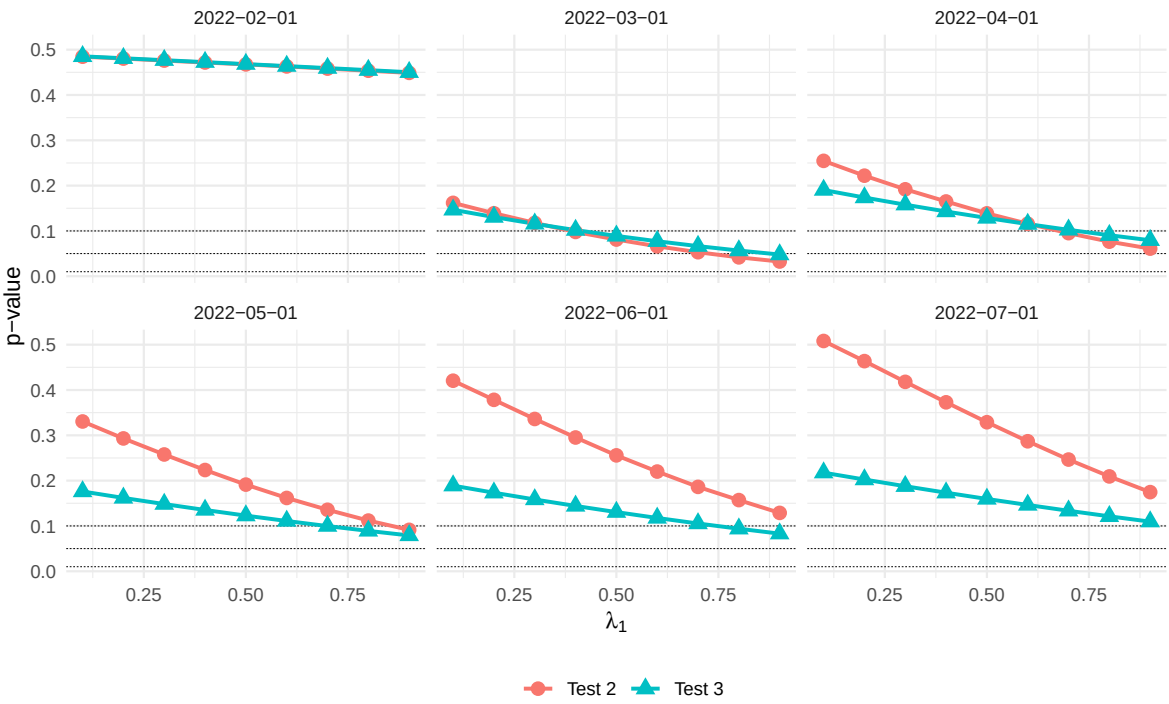


Figure 11: p-values for hypothesis tests as strength of monetary policy shock varies: Post-2021 residual covariance

Figure shows p-values for hypothesis tests described in Section 5.3, under a series of nulls, indexed by λ_1. This version uses the post-2021 residual covariance for $\hat{\Omega}_y$. Horizontal lines show 1, 5, and 10 percent confidence levels. The interpretation of λ_1 is that it represents a null where monetary policy is only λ_1 as effective as pre-COVID. Each panel computes the corresponding hypothesis tests on a window starting in February 2022 and ending in the titular month. The values for both lines when $\lambda_1 = 1$ thus match the corresponding columns in Table 6.

in the VAR. Every series shows an increase in the standard deviation, in some cases more than doubling. And so in Figure 11, we repeat the exercise in Figure 9, but replacing $\hat{\Omega}_y$ with the covariance matrix of the VAR residuals from 2021 onward. Given we use the same D_m matrix, the interpretation is that we are testing whether the impact of monetary policy has changed, allowing for any excess volatility in the data to be explained by changing variances of the other shocks. The results do somewhat mitigate our earlier findings on the transmission of monetary policy, suggesting that the evidence of given reductions in the impact of monetary policy are a little less strong. Nevertheless, the implication that policy was weaker than expected in at least March and April of 2022 remains robust. Moreover, this experiment sets a relatively high threshold for concluding that the impact of monetary policy has changed, since the residual shock covariance itself will be imprecisely estimated and including the full post-2021 period may not be representative of the population covariance during 2022.

6.4 Implications from a medium-scale macro model

In Section 6.2 we attempted to quantify how much monetary policy transmission might have weakened, considering the case where weaker monetary policy affects all variables equally. Although a reasonable starting point, this is somewhat restrictive. Even the very simplest macroeconomic models contradict this assumption.

For example, in the well-known standard New Keynesian model, both a steeper Phillips curve or a flatter intertemporal IS curve would weaken monetary transmission, in the sense of dampening the response of output to a monetary policy shock. However, the change in the impact of monetary policy on output varies across these two cases, and in neither scenario is it the same as the marginal impact on inflation. We show this explicitly for the textbook model of Galí (2015) in Appendix D. The intuition is that a change in the slope of the IS curve affects how a given change in the real rate affects growth of the output gap, leaving the relationship between the output gap and inflation unaffected. However, a change in the slope of the Phillips curve affects the relationship between the output gap and inflation. The result is that although a flatter intertemporal IS curve will always reduce the impact of a given monetary tightening on inflation and output, a steeper Phillips curve may result in a smaller output response but a *larger* inflation response to monetary tightening.[17]

Given this concern, we extend our framework for testing whether the data is consistent with not just weaker transmission in general, but try to identify the underlying economic source of that weakness. Specifically, we make the following assumptions:

1. A structural model is a mapping from a history of shocks $\epsilon_0, \ldots, \epsilon_t$ to data y_t parameterized by a vector θ.

2. Shocks are normal, independent of each other and their lags, and have unit variance.

3. The model is invertible, in that the solution can be approximated arbitrarily precisely by sufficiently long finite-order VAR:

$$y_t = \sum_{j=1}^{J} B_j(\theta) y_{t-j} + D(\theta)\epsilon_t$$

4. The first entry of the shock vector is a monetary shock ϵ_t^m, and the first column of $D(\theta)$ is the contemporaneous response of the model to this shock, which we write as $D_m(\theta)$.

5. For some baseline value of the parameters θ_0, the model produces the same monetary impulse responses as our empirical estimates. That is:

$$D_m(\theta_0) = \hat{D}_m \qquad\qquad B_j(\theta_0) = \hat{B}_j \ \forall j$$

[17] Under some conditions. Most notably, if monetary shocks are sufficiently persistent.

6. For the same θ_0, the model produces the same residual covariance as our empirical estimates:

$$D(\theta_0)D(\theta_0)' = \hat{\Omega}_y$$

If the preceding assumptions hold true, then when $\theta = \theta_0$ the model represents exactly the same view of monetary transmission (since the impulse responses are identical). Moreover, for any given observation y_t, the filter distribution in equation (7) is exactly the same whether one uses the empirical estimates or their model equivalents.[18] In other words, the structural model is an equivalent to our estimated statistical model in that it will rationalize a given set of observations as being consistent by exactly the same distribution set of monetary shocks.

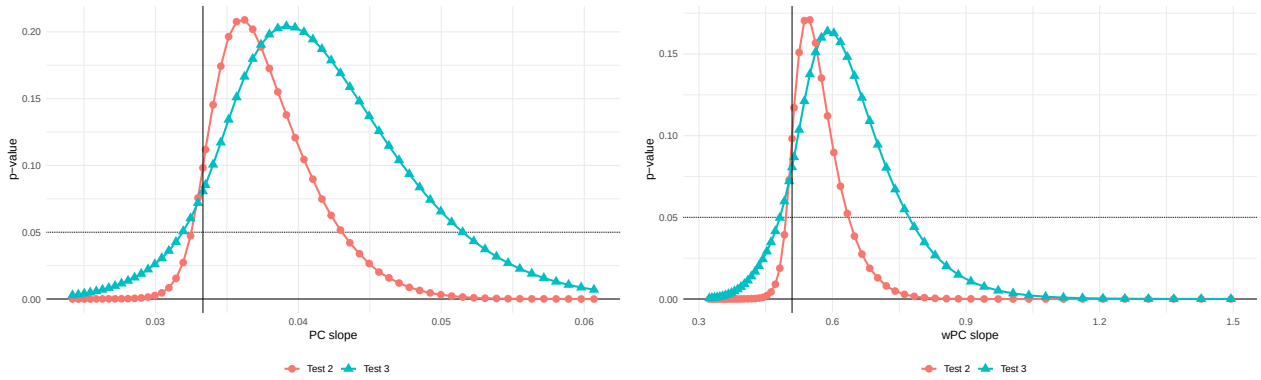
The advantage of using a structural model to rationalize our empirical results is that we can run more informative counterfactual tests rather than just uniformly weaker monetary transmission. In particular, for any alternative parameter value θ_1 we can re-run our hypothesis tests replacing $\hat{D}_M$ with $D_m(\theta_1)$, $\hat{B}_j$ with $B_j(\theta_1)$, and $\hat{\Omega}_y$ with $D(\theta_1)D(\theta_1)'$. This then re-formulates the null as a data generating process consistent with a structural model with the new parameters. The alternate parameter values required to fail to reject the null can then offer some insight into what might have changed to make the transmission of monetary policy weaker in 2022.

To implement this in practice, we take the well-known model of Smets and Wouters (2007) as a baseline for the US economy. In our main exercise, we parametrize the model to fit the estimated impulse response. This determines the values of the vector of parameters θ_0 (see Appendix D.2 for details).[19] We then consider alternative values for 3 elements of the parameter vector: the Calvo reset probability for prices, the Calvo reset probability for wages, and the inter-temporal elasticity of substitution. These have different economic impacts, affecting the wage Phillips curve, the price Phillips curve, and the IS curve respectively. For each alternate parameter vector θ_1, we re-solve the model and compute the response to a monetary shock. This gives values for $D_m(\theta_1)$, $B_j(\theta_1)$, and $D(\theta_1)$, which we can use to calculate the distribution of the test statistic under the alternate null.

Figure 12 reports the results of this exercise for alternate values of the two Calvo reset parameters. This shows that perturbing the Calvo probabilities so that the slope of either the wage or price Phillips curves increase relative to the baseline raises the p-value of the test. That is, the data in 2022 cannot reject the possibility that the wage or price Phillips curves steepened. Of course, this increase is not limitless; for sufficiently steep Phillips curves

[18] Why? Because $B_j = B_j(\theta_0)$ means that the reduced-form residual $\hat{\epsilon}_t^y$ is the same, and because $D_m(\theta_0) = \hat{D}_m$ and $D(\theta_0)D(\theta_0)' = \hat{\Omega}_y$ means that the other ingredients of this formula are identical.

[19] The model is not able to match our impulse responses very closely. So to check the extent to which this might affect our results, we also apply our method using the original model calibration of Smets and Wouters (2007). The results are similar, see Appendix B.6.



(a) Price Phillips Curve

(b) Wage Phillips Curve

Figure 12: p-values for alternate structural parameters

Figures show p-values for tests of alternate nulls using the data generating process from the Smets and Wouters (2007) model. For each point in panel 12a we recompute the p-value of the tests on the February-July 2022 data using the values of $D_m(\theta_1)$, $B_j(\theta_1)$, and $D(\theta_1)$ implied by the model solution when the price Calvo reset probability is varied. We then plot this (on the y-axis) against the Phillips curve (on the x-axis). Panel 12b does the same but for the wage Calvo reset probability. The horizontal dashed line shows the 5 percent critical value and the vertical line the baseline value, matching our empirical exercise.

the p-value falls again; very steep Phillips curves are also inconsistent with the 2022 data.

To capture the range of structural interpretations of our model consistent with the data, Table 7 reports the range of slope coefficient with p-values in excess of 5 percent, in the columns labeled "acceptance region". These correspond to the parts of Figure 12 above the 5 percent horizontal line. In relative terms the acceptance regions for the two Phillips curves are remarkably similar. The data are consistent with both price and wage Phillips curves which are as much as 25 to 50 percent larger than the baseline under tests 2 and 3 respectively.[20] However, the results from altering the IS curve are somewhat counterintuitive, implying that the data would be consistent with a steeper IS curve. One important difference here is that the relative change in the IS curve is much smaller, implying that the data would only admit a fractionally steeper IS curve before rejecting the model once more. Moreover, to the extent that p-values can be treated as a metric of how consistent data are with a given hypothesis, the higher peak p-values of the alternate price Phillips curve models suggest that this is a more plausible interpretation of our results.[21]

Overall, we consider the results of this section as being broadly indicative of the idea that the weaker transmission of monetary policy was due to a steepening of the price (or wage) Phillips curve. A steeper IS curve could be an alternate explanation, but the magnitude would be much smaller and the result less statistically plausible. However, one should not over-interpret these results. We think of them as offering a plausible approximate structural

[20] Because $0.042/0.033 \simeq 0.62/0.51 \simeq 1.25$ and $0.051/0.033 \simeq 0.75/0.51 \simeq 1.5$

[21] Re-running this exercise using the Smets and Wouters (2007) original calibration produces very similar results, see Table A2

interpretation for our empirical findings, rather than conclusive evidence.

Slope of	Baseline	Parameter changed	Test 2		Test 3	
			Acceptance region	Max p-value	Acceptance region	Max p-value
Phillips Curve	0.033	ξ_p, Price Calvo probability	(0.033, 0.043)	0.21	(0.032, 0.051)	0.2
Wage Phillips Curve	0.51	ξ_w, Wage Calvo probability	(0.5, 0.63)	0.17	(0.49, 0.76)	0.16
IS Curve	1.7	σ_c, Inter-termporal EoS	(1.7, 1.8)	0.16	(1.7, 2)	0.15

Table 7: Summary results: changing IS and Phillips curve slopes in a structural model, impulse-response matched parameter values.

Table reports the results of perturbing the Smets and Wouters (2007) model of the US economy by various parameters and using the corresponding impulse response functions in our tests. The model is estimated using the impulse response matching approach, see Appendix D.2.

7 Conclusions

In this paper we propose a general method for assessing whether a dynamic data generating process has changed when one has externally-identified shocks. We applied it to the question of whether the transmission of monetary policy in the US was different during the 2022 tightening cycle. We conclude that during 2022 the macro data and the shocks were not consistent with unchanged monetary transmission and that monetary policy was probably around 25 percent less effective than would have been expected pre-COVID. When we allow for changing shock variance, our findings are weakened but not overturned. When our series also includes policymakers' speeches, they are strengthened.

The core of this paper is in the method for inverting the VAR to back out specific shocks. The strength of our approach is twofold. First, it takes seriously the concern that other shocks may be acting to mask the impact of monetary policy, and thus giving the appearance of an altered transmission mechanism. Second, by maintaining a fixed pre-test null and by focusing on the changing response of just monetary shocks, we can get reasonable statistical power even when looking at relatively short periods. This contrasts to other approaches, such as time-varying parameter methods or subsample re-estimation, where the quantity of changing parameters impedes inference.

This generality come at a price, though. Although we can robustly test whether monetary transmission changed, it is rather harder to say why. We attempt to offer some explanation for our results in Section 6 by coercing a standard macro model to match our pre-COVID impulse responses and seeing how much it must be changed to fail to reject the 2022 data. And although this gives us some evidence that a steeper Phillips curve may be to blame, we do not see this as conclusive.

More generally, though, our method has broad scope for wider application, to almost any setting where one has a well-identified shock and wishes to assess whether the propagation of that shock has changed. For example, using narrative identification of fiscal shocks one

might be able to more precisely examine the extent to which public spending multipliers are state-dependent. Similarly, an identified set of exchange rate shocks could allow one to assess whether the strength of the pass-through to inflation has changed in response to a new policy environment. We leave these applications for future research.

8 Acknowledgements

We thank Nigel Chalk, Mai Chi Dao, Jesper Lindé, Juan Passadore, Linda Tesar, and Ken West for helpful comments, and Seminar Participants at the IMF's Western Hemisphere Departmental Seminar and Monetary Policy Advisory Group, and attendees at the 2024 LACE/LAMES conference. The views expressed herein are those of the authors and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

References

Ascari, Guido and Timo Haber, "Non-linearities, state-dependent prices and the transmission mechanism of monetary policy," *The Economic Journal*, 2022, *132* (641), 37–57.

Ball, Laurence, Daniel Leigh, and Prachi Mishra, "Understanding US Inflation during the COVID-19 Era," *Brookings Papers on Economic Activity*, 2022, *53* (2 (Fall)), 1–80.

Bauer, Michael D., Carolin E. Pflueger, and Adi Sunderam, "Changing Perceptions and Post-Pandemic Monetary Policy," 2024. Proceedings, *Jackson Hole Economic Symposium*.

Berger, David, Konstantin Milbradt, Fabrice Tourre, and Joseph Vavra, "Mortgage prepayment and path-dependent effects of monetary policy," *American Economic Review*, 2021, *111* (9), 2829–2878.

Bernstein, Joshua, "A model of state-dependent monetary policy," *Journal of Monetary Economics*, 2021, *117*, 904–917.

Casella, George and Roger Berger, *Statistical inference*, CRC Press, 2024.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, February 2005, *113* (1), 1–45.

Cohen, Elior, "Post-Pandemic Labor Shortages Have Limited the Effect of Monetary Policy on the Labor Market," *Kansas City Fed Economic Bulletin*, 2023.

D'Amico, Stefania and Thomas King, "Past and Future Effects of the Recent Monetary Policy Tightening," *Chicago Fed Letter*, 2023, (483).

Eichenbaum, Martin, Sergio Rebelo, and Arlene Wong, "State-dependent effects of monetary policy: The refinancing channel," *American Economic Review*, 2022, *112* (3), 721–761.

English, Bill, Kristin Forbes, and Angel Ubide, *Monetary Policy Responses to the Post-Pandemic Inflation*, CEPR Press, 2024.

Fosso, Luca, Ine V. Robays, and Kasper Goosen, "Transmission of US monetary policy to the labour market: is this time different?," *SUERF Policy Brief*, December 2023, (759).

Galí, Jordi, *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*, Princeton University Press, 2015.

Gertler, Mark and Peter Karadi, "Monetary policy surprises, credit costs, and economic activity," *American Economic Journal: Macroeconomics*, 2015, *7* (1), 44–76.

Glover, Andrew and Johnson Oliyide, "Why Haven't Recent Rate Increases Slowed the Economy More? Look to Unusually Low Private-Lending Spreads," *Economic Bulletin*, August 2024, pp. 1–4.

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

_ , Hatice Gökçe Karasoy-can, and Sang Seok Lee, "Stock Market's Assessment of Monetary Policy Transmission: The Cash Flow Effect," *Journal of Finance*, 2022, *77* (4), 2375–2421.

Jarociński, Marek, "Central bank information effects and transatlantic spillovers," *Journal of International Economics*, 2022, *139*, 103683.

Jarociński, Marek and Peter Karadi, "Deconstructing monetary policy surprises—the role of information shocks," *American Economic Journal: Macroeconomics*, 2020, *12* (2), 1–43.

Kuttner, Kenneth N, "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market," *Journal of Monetary Economics*, 2001, *47* (3), 523–544.

Mertens, Karel and Morten O Ravn, "The dynamic effects of personal and corporate income tax changes in the United States," *American Economic Review*, 2013, *103* (4), 1212–1247.

Olea, José L Montiel, James H Stock, and Mark W Watson, "Inference in structural vector autoregressions identified with an external instrument," *Journal of Econometrics*, 2021, *225* (1), 74–87.

Paul, Pascal, "The time-varying effect of monetary policy on asset prices," *Review of Economics and Statistics*, 2020, *102* (4), 690–704.

Smets, Frank and Rafael Wouters, "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, June 2007, *97* (3), 586–606.

Stedman, Karlye Dilts and Emily Pollard, "Why Has Monetary Policy Tightening Not Cooled the Labor Market Enough to Quell Inflation?," *Economic Bulletin*, 2023, pp. 1–4.

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

Swanson, Eric and Vishuddhi Jayawickrema, "Speeches by the fed chair are more important than FOMC announcements: An improved high-frequency measure of us monetary policy shocks," *Working paper.*, 2023.

Swanson, Eric T., "The Importance of Fed Chair Speeches as a Monetary Policy Tool," *AEA Papers and Proceedings*, May 2023, *113*, 394–400.

Tenreyro, Silvana and Gregory Thwaites, "Pushing on a string: US monetary policy is less powerful in recessions," *American Economic Journal: Macroeconomics*, 2016, *8* (4), 43–74.

Waggoner, Daniel F and Tao Zha, "Conditional forecasts in dynamic multivariate models," *Review of Economics and Statistics*, 1999, *81* (4), 639–651.

Online Appendix: Not for Publication

A Data

A.1 Details on the baseline shock series

Let $p_{t_1}^{FF4}$ be the last traded price of the fourth Federal Funds Future (FF4) before the event window starts, and $p_{t_2}^{FF4}$ be the first traded price of FF4 once the event window ends. The raw shock for FF4 ξ_t^{FF4} is calculated as

$$\xi_t^{FF4} = p_{t_2}^{FF4} - p_{t_1}^{FF4}$$

Let $p_{t_1}^{SP500}$ be the last traded value of the S&P 500 index before the event window starts, and $p_{t_2}^{SP500}$ be the first value of the S&P 500 index once the event window ends. The raw shock for S&P 500 ξ_t^{SP500} is calculated as

$$\xi_t^{SP500} = \ln p_{t_2}^{SP500} - \ln p_{t_1}^{SP500}$$

A.2 Speeches

The fundamental idea to construct high-frequency monetary shocks from speeches by members of the FOMC is the same as for FOMC meetings. We take the change in asset prices around the event dates, as described in section 3.1 and appendix A.1 . To follow this procedure, we have to collect exact dates and times of the respective speeches, and choose an appropriate event time window. We collect dates of speeches from the homepage of the Federal Reserve Board. Starting from year 2010, this page also links the speech as a pdf document, including the time the speech was released (including the timezone). We follow Swanson and Jayawickrema (2023) and set the event window to 120 minutes. Using this procedure, we are able to use 714 speeches ranging from 2010 to 2023 to construct raw shocks from both the fourth Federal Funds Futures (FF4) and the S&P 500 index. As with regular meetings, we aggregate to monthly frequency by taking the sum of the individual shocks

within a given month.

A.3 Extra time series plots

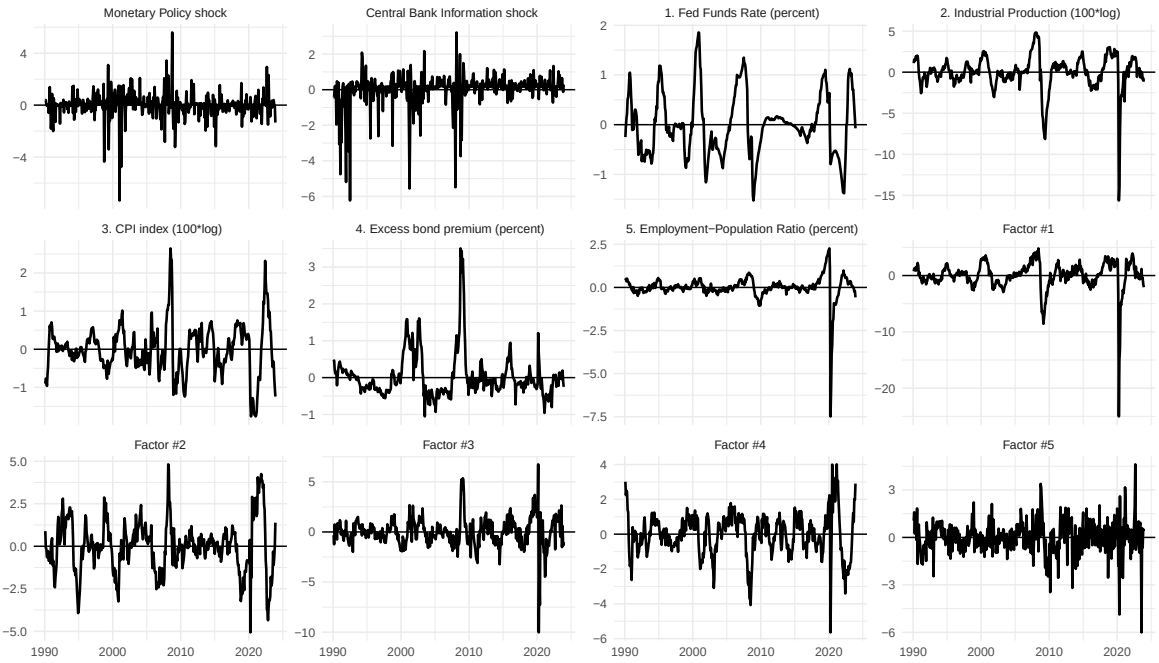


Figure A1: Data in the baseline VAR

Figure shows the data series in the baseline vector autogression including factors, after transformations.

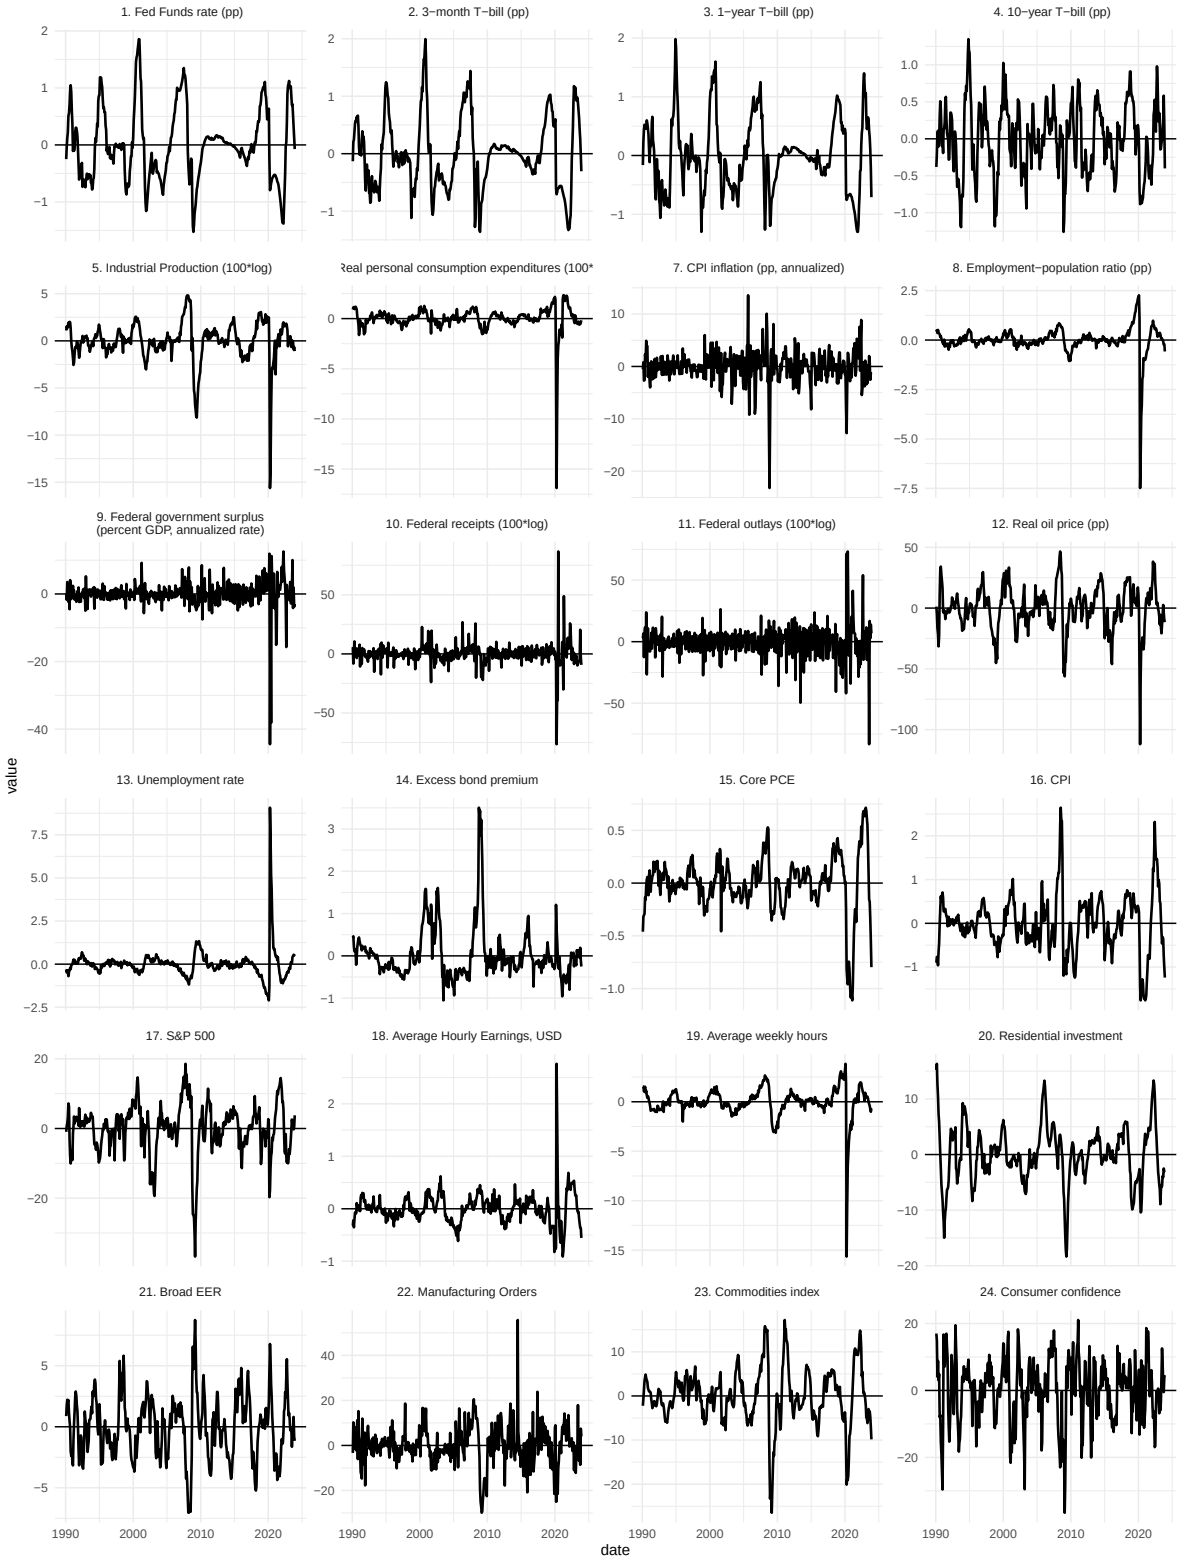


Figure A2: All time series data

Figure shows the full set of data in Table 2, after transformations.

B Additional results

B.1 More results under the baseline

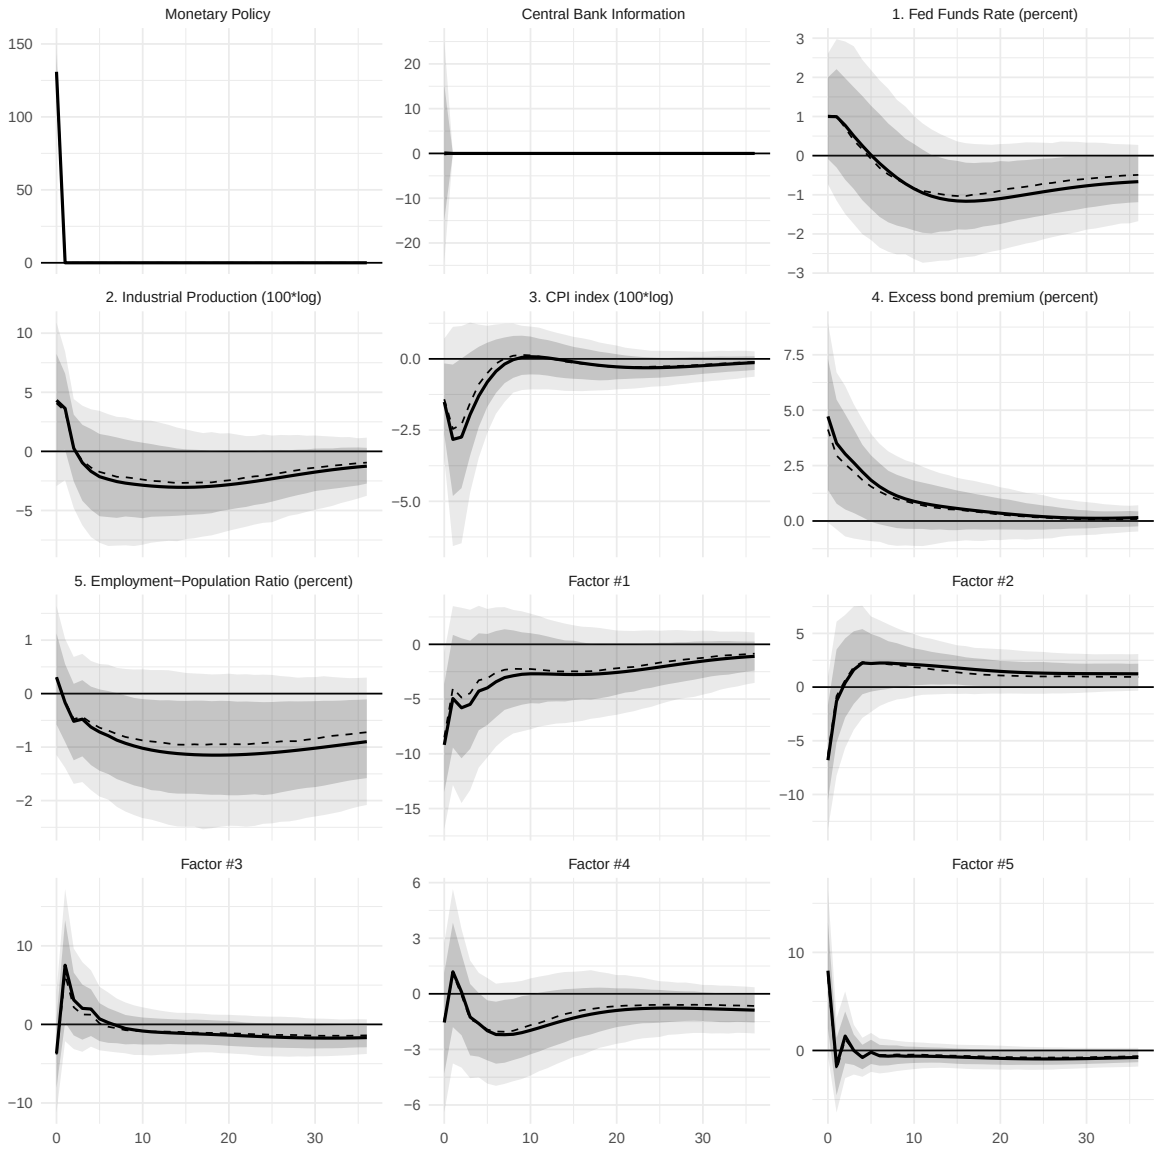


Figure A3: Impulse responses: Monetary policy shock

Figure shows the response of the headline variables to a monetary policy shock, scaled to a 1 percentage point increase in the Federal Funds Rate. Solid lines are point estimates and shaded regions are the 68 and 90 percent confidence intervals from a bootstrap with $K = 1000$ replications. Dashed lines show median responses from the bootstrap.

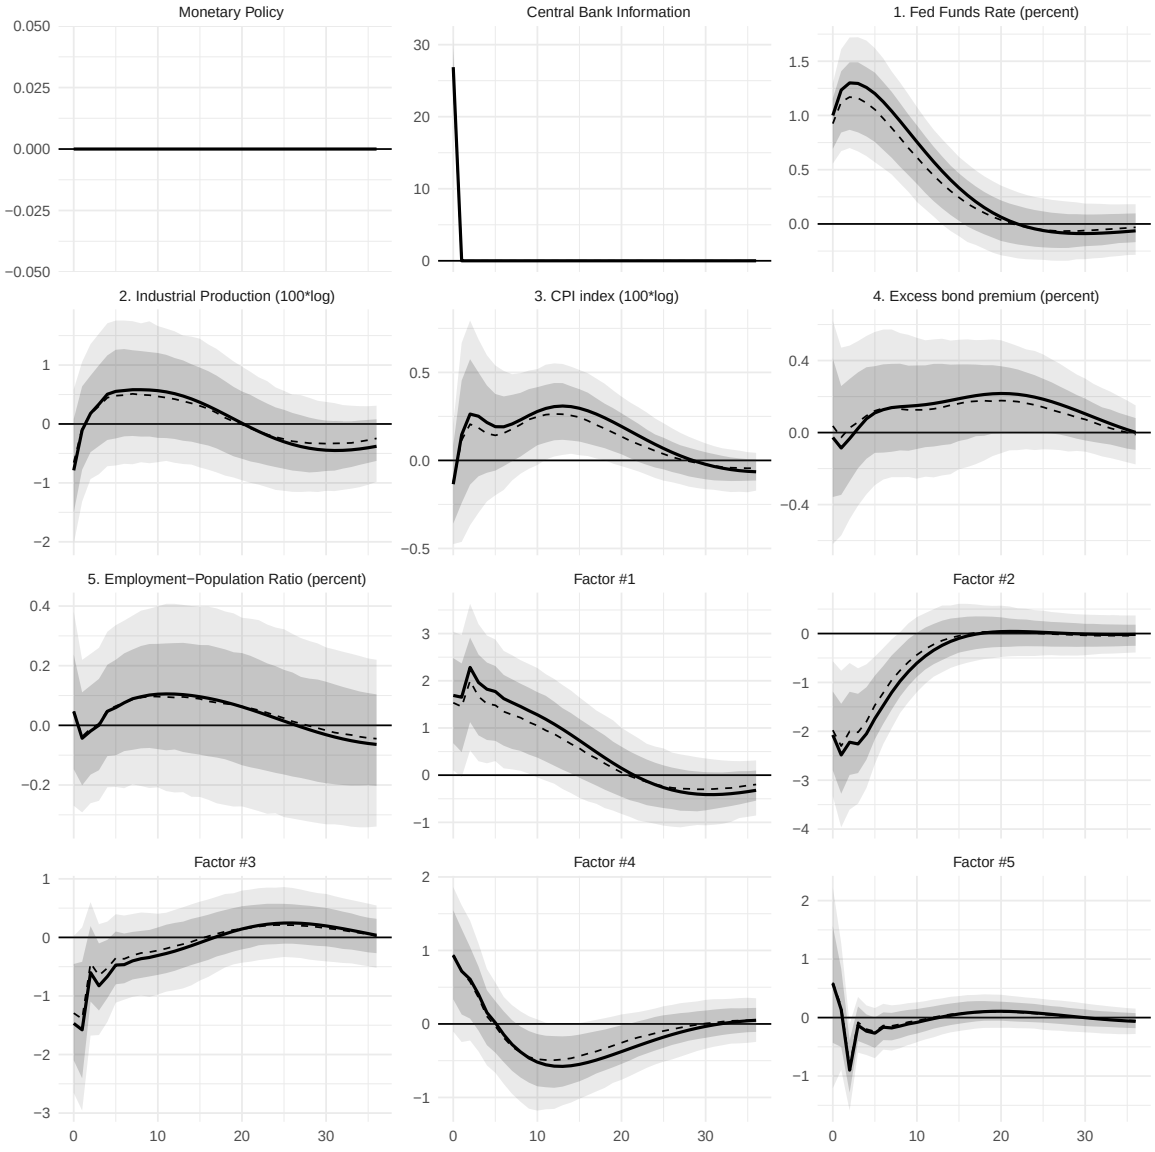


Figure A4: Impulse responses: Central Bank information shock

Figure shows the response of the headline variables to a monetary policy shock, scaled to a 1 percentage point increase in the Federal Funds Rate. Solid lines are point estimates and shaded regions are the 68 and 90 percent confidence intervals from a bootstrap with $K = 1000$ replications. Dashed lines show median responses from the bootstrap.

B.2 Alternative specifications

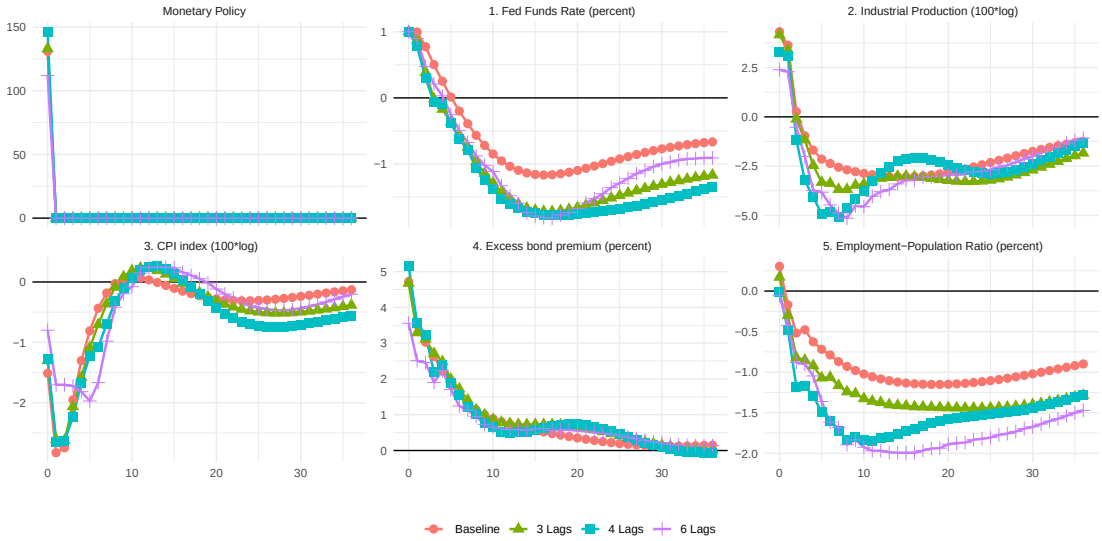


Figure A5: Impulse responses: Monetary policy shock, different lag structures

Figure shows the point estimates of the response of a factor-augmented VAR to a monetary policy shock, scaled to a 1 percentage point increase in the interest rate. Lines differ in their lag length, which in the baseline is two periods, chosen by AIC. All specifications include five factors, responses of which are not shown here.

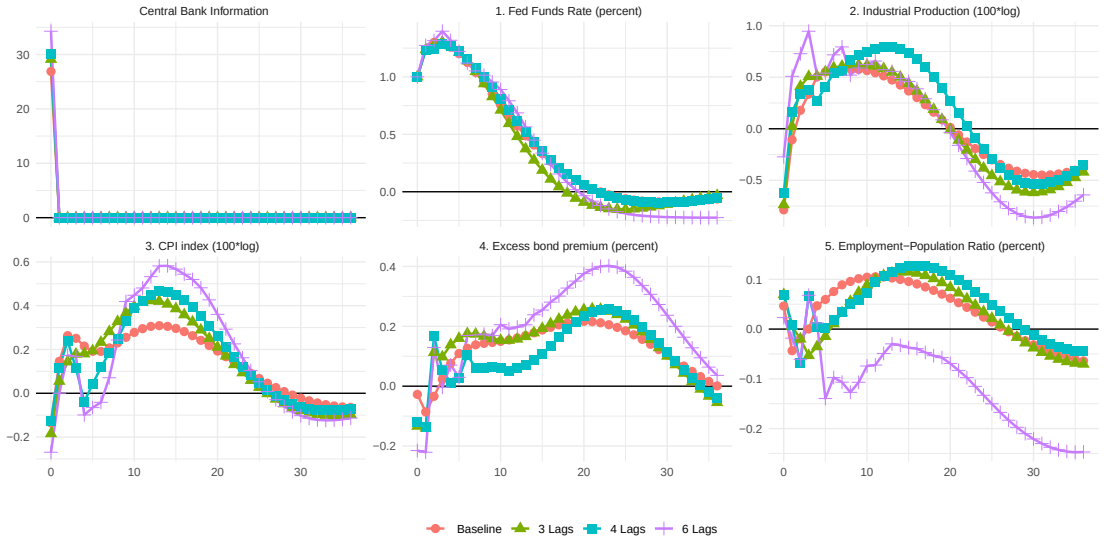


Figure A6: Impulse responses: Central Bank information shock, different lag structures

Figure shows the point estimates of the response of a factor-augmented VAR to a Central Bank information shock, scaled to a 1 percentage point increase in the interest rate. Lines differ in their lag length, which in the baseline is two periods, chosen by AIC. All specifications include five factors, responses of which are not shown here.

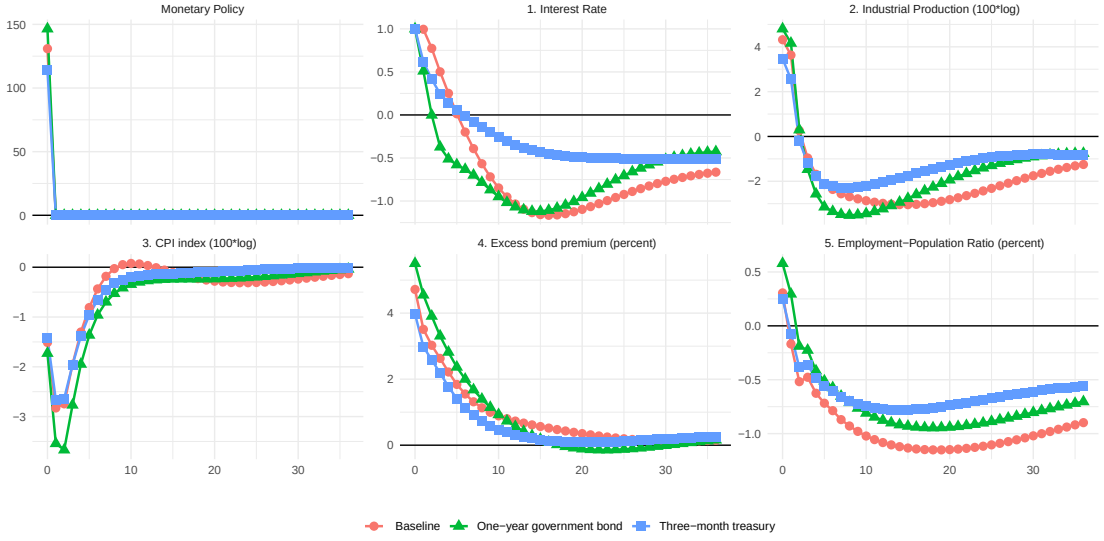


Figure A7: Impulse responses: Monetary policy shock, different interest rate series

Figure shows the point estimates of the response of a factor-augmented VAR to a monetary policy shock, scaled to a 1 percentage point increase in the interest rate. Lines differ in their measure of the interest rate, which in the baseline is the Federal Funds Rate. All specifications include five factors, responses of which are not shown here.

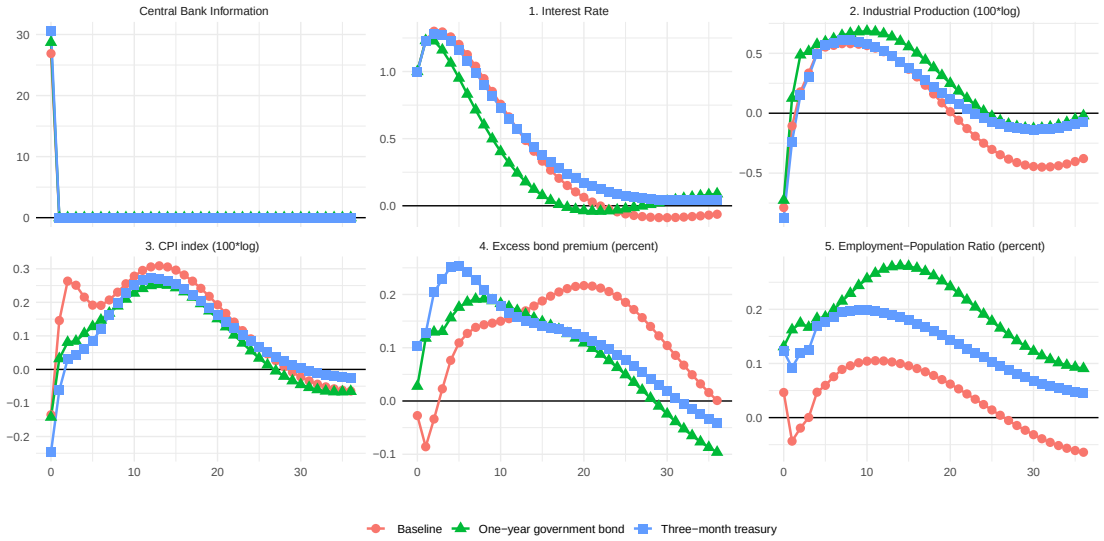


Figure A8: Impulse responses: Central Bank information shock, different interest rate series

Figure shows the point estimates of the response of a factor-augmented VAR to a Central Bank information shock, scaled to a 1 percentage point increase in the interest rate. Lines differ in their measure of the interest rate, which in the baseline is the Federal Funds Rate. All specifications include five factors, responses of which are not shown here.

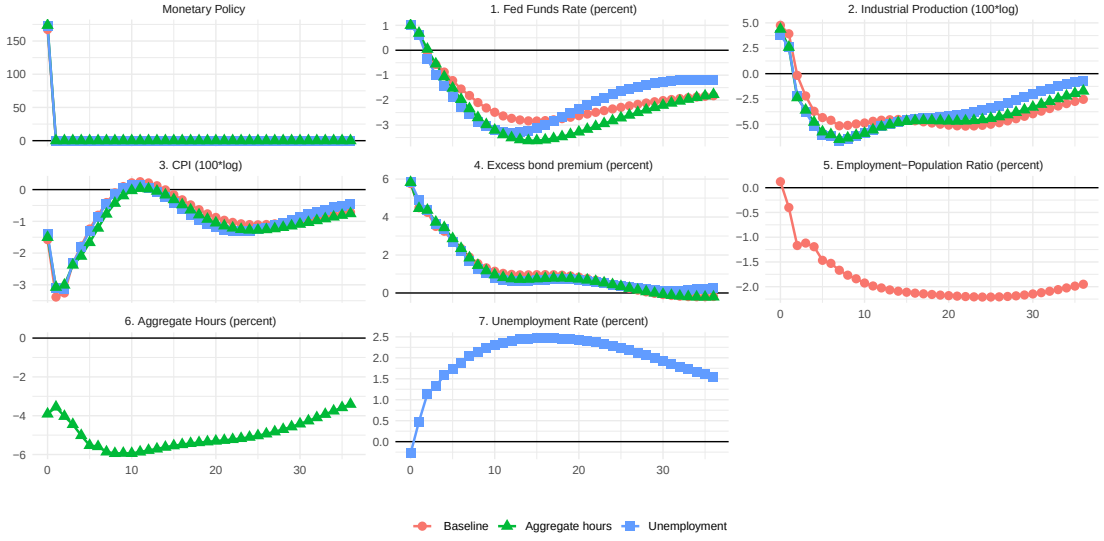


Figure A9: Impulse responses: Monetary policy shock, different labor series

Figure shows the point estimates of the response of a factor-augmented VAR to a monetary policy shock, scaled to a 1 percentage point increase in the interest rate. Lines differ in their labor market variable, which in the baseline is the employment-population ratio. All specifications include five factors, responses of which are not shown here.

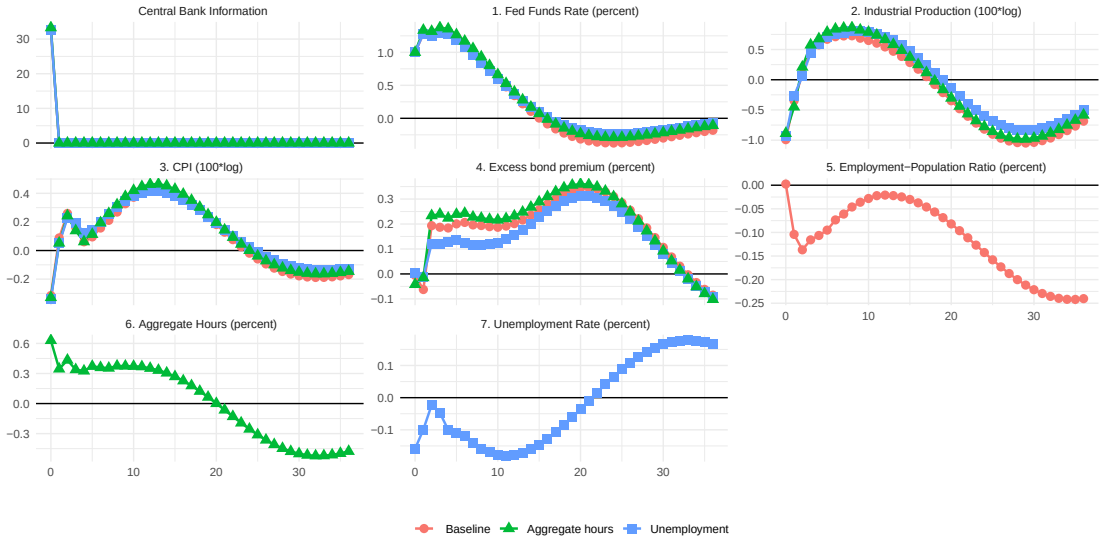


Figure A10: Impulse responses: Central Bank information shock, different labor series

Figure shows the point estimates of the response of a factor-augmented VAR to a monetary policy shock, scaled to a 1 percentage point increase in the interest rate. Lines differ in their labor market variable, which in the baseline is the employment-population ratio. All specifications include five factors, responses of which are not shown here.

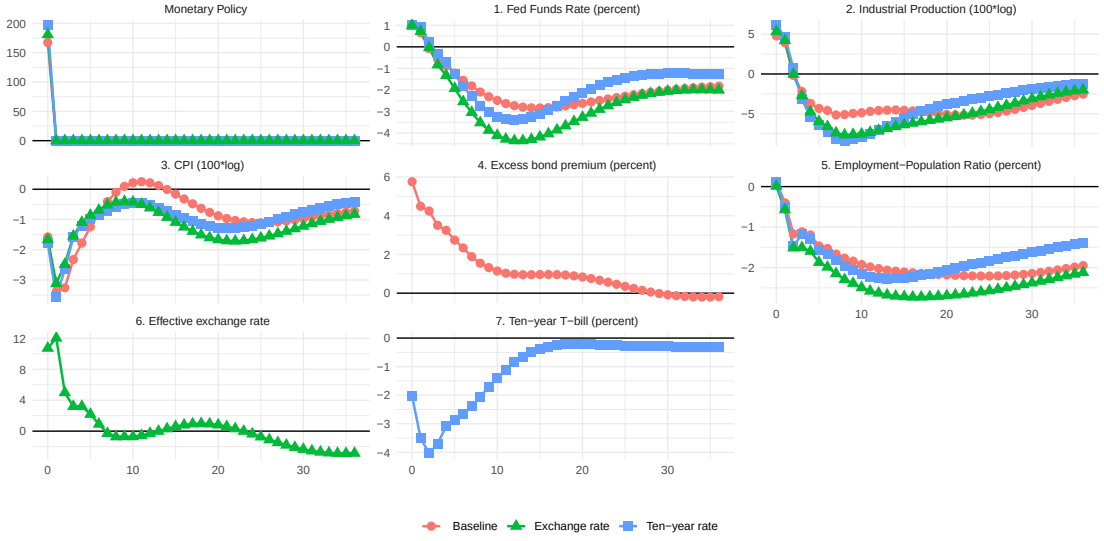


Figure A11: Impulse responses: Monetary policy shock, different financial series

Figure shows the point estimates of the response of a factor-augmented VAR to a monetary policy shock, scaled to a 1 percentage point increase in the interest rate. Lines differ in their financial variable, which in the baseline is the excess bond premium. All specifications include five factors, responses of which are not shown here.

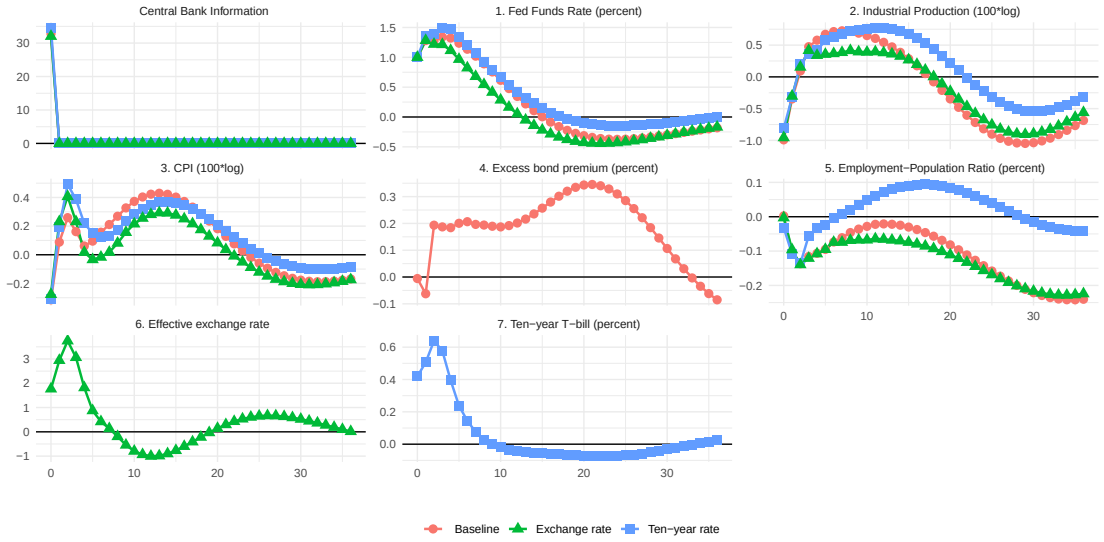


Figure A12: Impulse responses: Central Bank information shock, different financial series

Figure shows the point estimates of the response of a factor-augmented VAR to a monetary policy shock, scaled to a 1 percentage point increase in the interest rate. Lines differ in their financial variable, which in the baseline is the excess bond premium. All specifications include five factors, responses of which are not shown here.

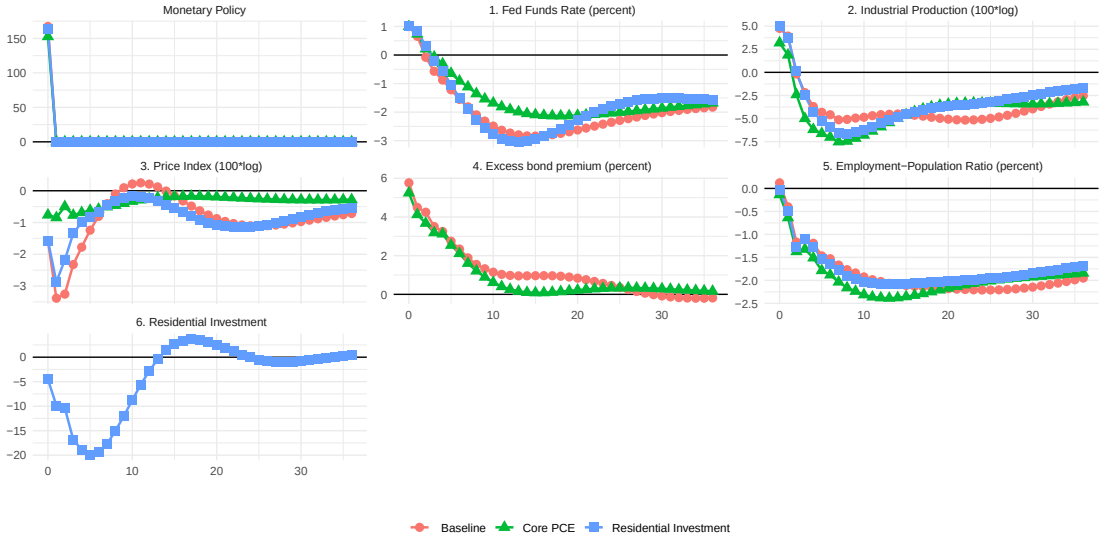


Figure A13: Impulse responses: Monetary policy shock, different other variables

Figure shows the point estimates of the response of a factor-augmented VAR to a monetary policy shock, scaled to a 1 percentage point increase in the interest rate. Lines differ in the variables that they include. The "Core PCE" specification replaces the CPI index with core PCE and is shown in panel 3. The "Residential Investment" specification replaces the employment-population ratio with log residential investment. All specifications include five factors, responses of which are not shown here.

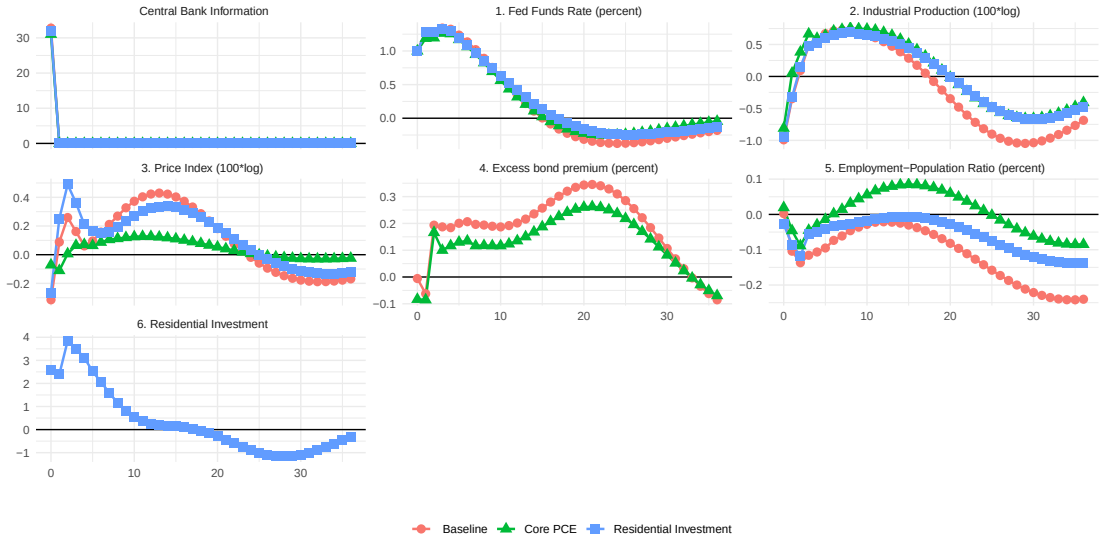


Figure A14: Impulse responses: Central Bank information shock, different other variables

Figure shows the point estimates of the response of a factor-augmented VAR to a monetary policy shock, scaled to a 1 percentage point increase in the interest rate. Lines differ in the variables that they include. The "Core PCE" specification replaces the CPI index with core PCE and is shown in panel 3. The "Residential Investment" specification replaces the employment-population ratio with log residential investment. All specifications include five factors, responses of which are not shown here.

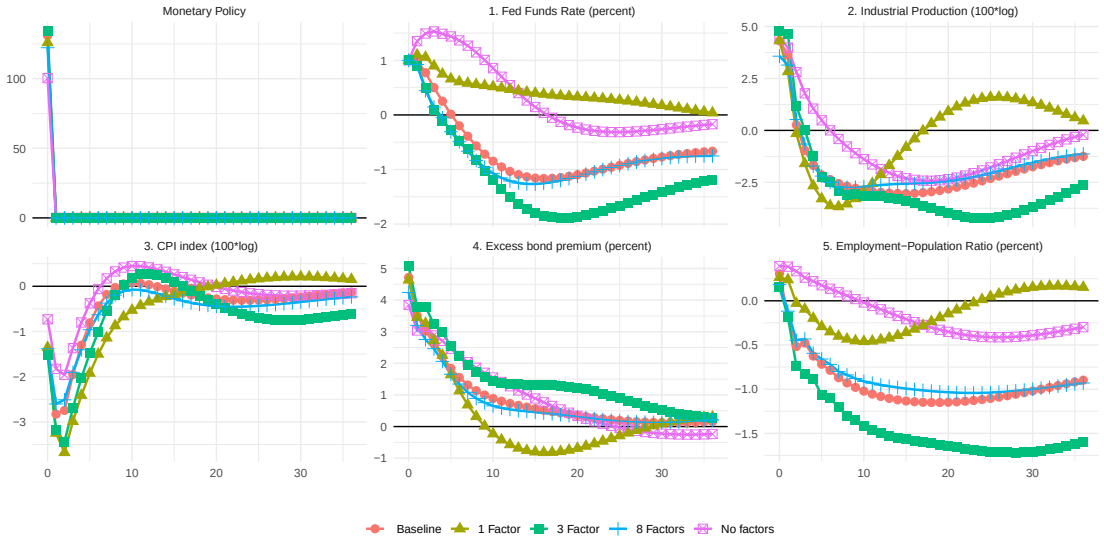


Figure A15: Impulse responses: Monetary policy shock, different number of factors

Figure shows the point estimates of the response of a factor-augmented VAR to a monetary policy shock, scaled to a 1 percentage point increase in the interest rate. Lines differ in the number of factors, which in the baseline is five. Responses of factors are are not shown here.

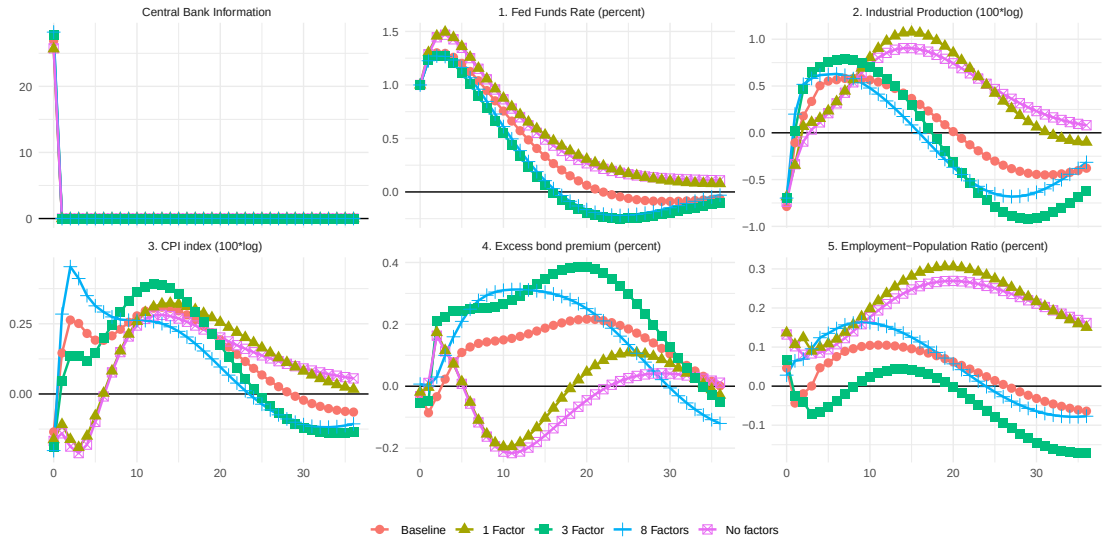


Figure A16: Impulse responses: Central Bank information shock, different number of factors

Figure shows the point estimates of the response of a factor-augmented VAR to a Central Bank information shock, scaled to a 1 percentage point increase in the interest rate. Lines differ in the number of factors, which in the baseline is five. Responses of factors are are not shown here.

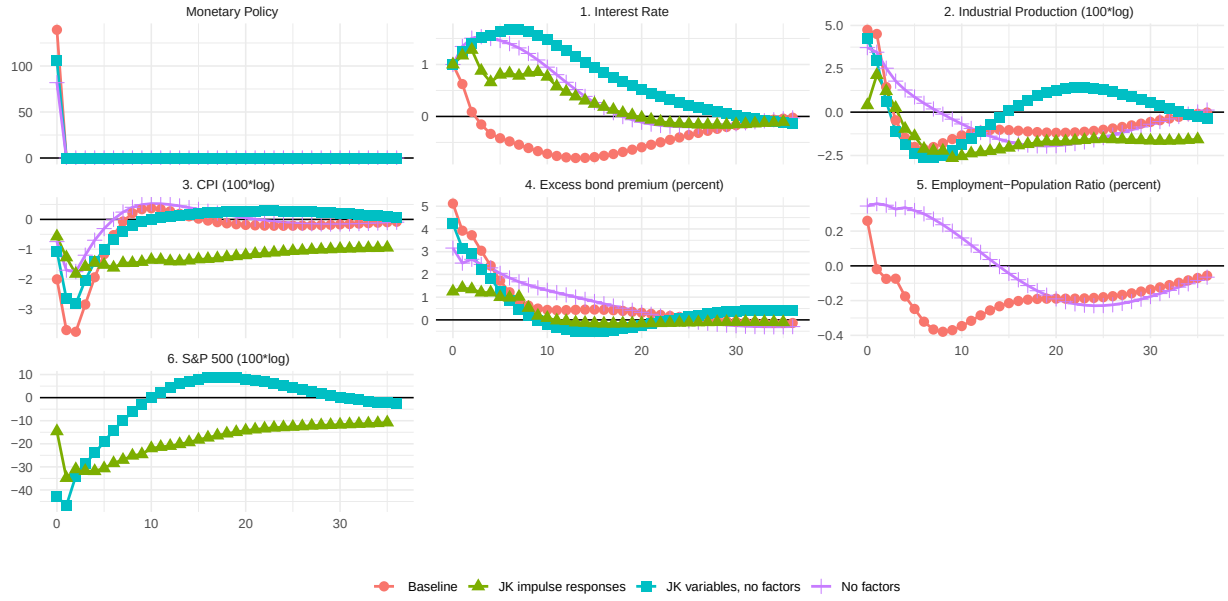


Figure A17: Impulse responses: Monetary policy shock, February 1992-December 2016.

Figure shows the point estimates of the response of a factor-augmented VAR to a monetary policy shock, scaled to a 1 percentage point increase in the interest rate. The "JK impulse responses" are the scaled impulse responses in Figure C3 of the Online Appendix of Jarociński and Karadi (2020) (which is monthly and so we use instead of their quarterly headline results). The "Baseline" line corresponds to our baseline specification, re-estimated on the matching sample: February 1992-December 2016. This specification retains factors, the responses of which are not shown here. The "No factors" line repeats our baseline estimation on this sample, just excluding factors. And "JK variables, no factors" matches the Jarociński and Karadi specification, excluding factors and replacing the Federal Funds Rate and employment-population ratio with the three-month treasury bill rate and S&P 500 index.

B.3 Hypothesis tests under alternative specifications

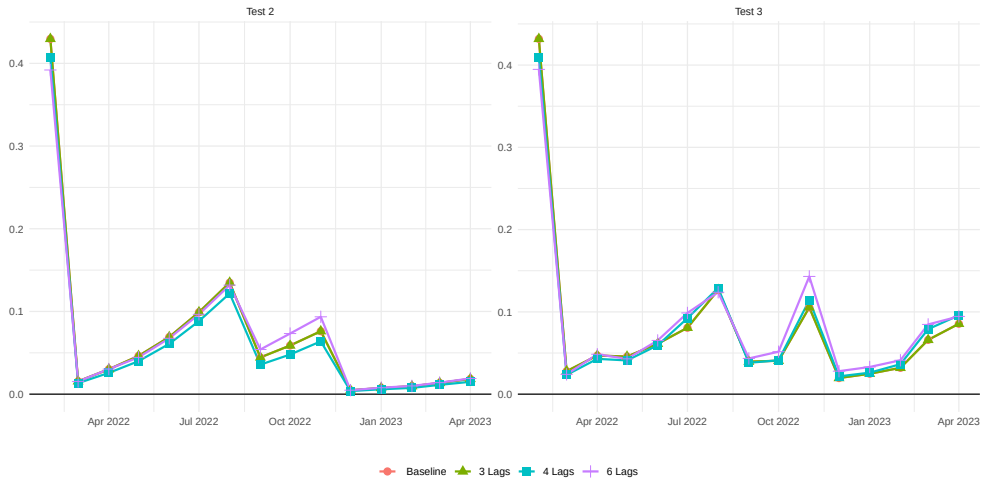


Figure A18: Hypothesis tests: Monetary policy shock, different lag structures

Figure shows the p-values of the rolling hypothesis tests that the impact of the monetary policy shock is unchanged in an expanding window starting in February 2022. Tests 2 and 3 test, respectively, whether the transmission of monetary policy has changed at some point and on average in the test window. Lines differ in their lag length, which in the baseline is two periods, chosen by AIC. All specifications also include five factors.

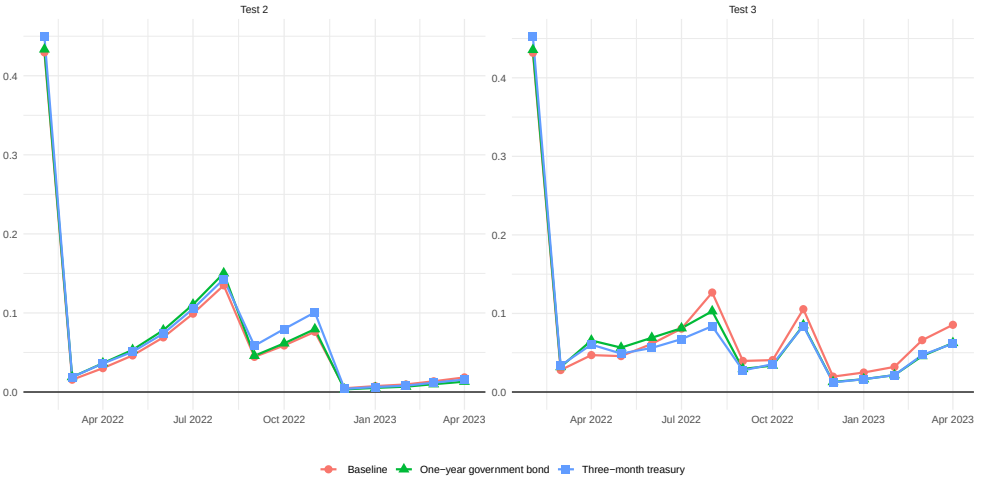


Figure A19: Hypothesis tests: Monetary policy shock, different interest rate series

Figure shows the p-values of the rolling hypothesis tests that the impact of the monetary policy shock is unchanged in an expanding window starting in February 2022. Tests 2 and 3 test, respectively, whether the transmission of monetary policy has changed at some point and on average in the test window. Lines differ in their measure of the interest rate, which in the baseline is the Federal Funds Rate. All specifications also include five factors.

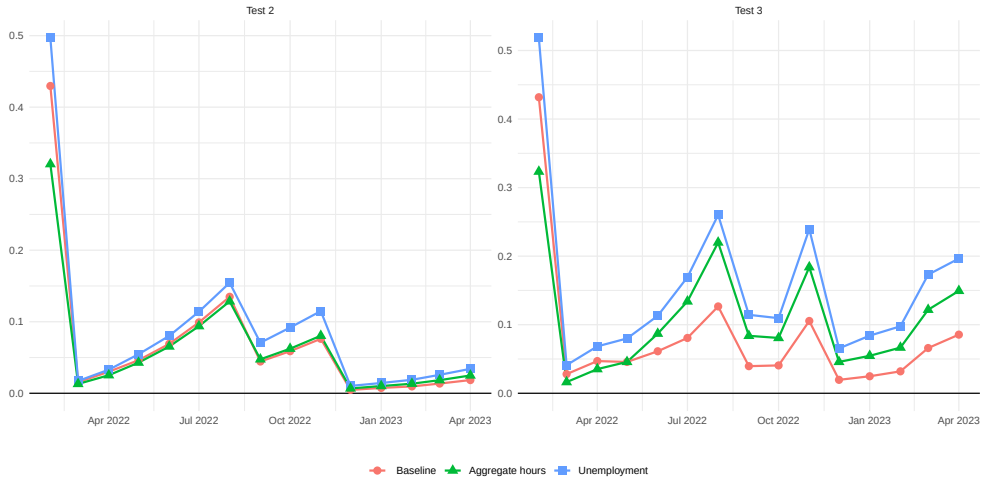


Figure A20: Hypothesis tests: Monetary policy shock, different labor series

Figure shows the p-values of the rolling hypothesis tests that the impact of the monetary policy shock is unchanged in an expanding window starting in February 2022. Tests 2 and 3 test, respectively, whether the transmission of monetary policy has changed at some point and on average in the test window. Lines differ in their labor market variable, which in the baseline is the employment-population ratio. All specifications also include five factors.

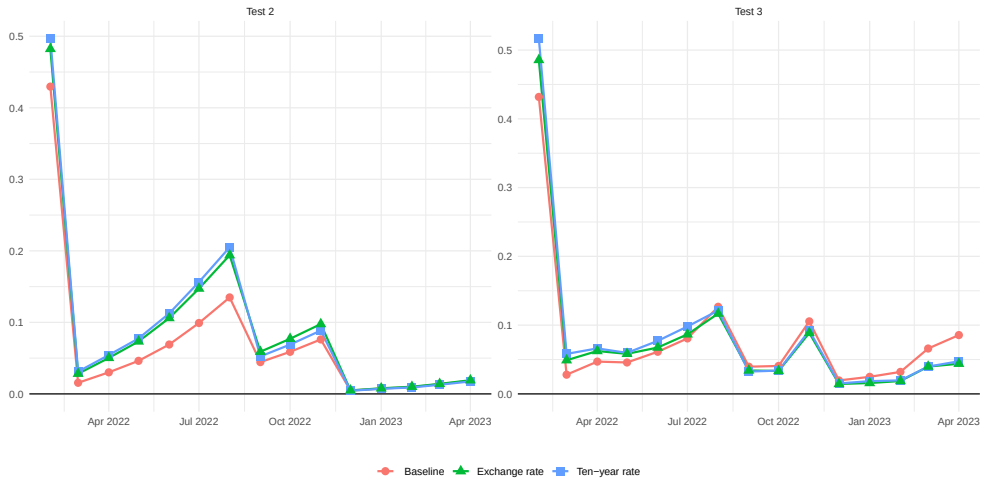


Figure A21: Hypothesis tests: Monetary policy shock, different financial series

Figure shows the p-values of the rolling hypothesis tests that the impact of the monetary policy shock is unchanged in an expanding window starting in February 2022. Tests 2 and 3 test, respectively, whether the transmission of monetary policy has changed at some point and on average in the test window. Lines differ in their financial variable, which in the baseline is the excess bond premium. All specifications also include five factors.

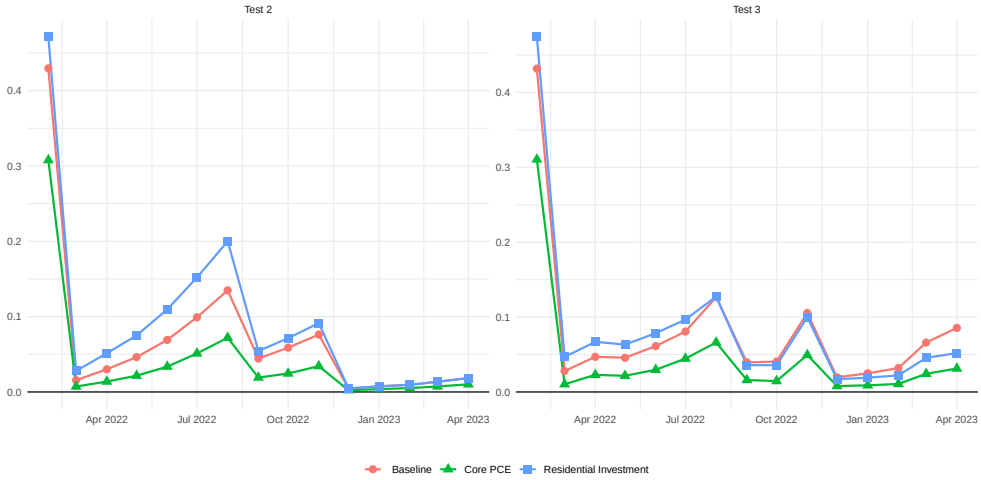


Figure A22: Hypothesis tests: Monetary policy shock, different other variables

Figure shows the p-values of the rolling hypothesis tests that the impact of the monetary policy shock is unchanged in an expanding window starting in February 2022. Tests 2 and 3 test, respectively, whether the transmission of monetary policy has changed at some point and on average in the test window. The "Core PCE" specification replaces the CPI index with core PCE and is shown in panel 3. The "Residential Investment" specification replaces the employment-population ratio with log residential investment. All specifications also include five factors.

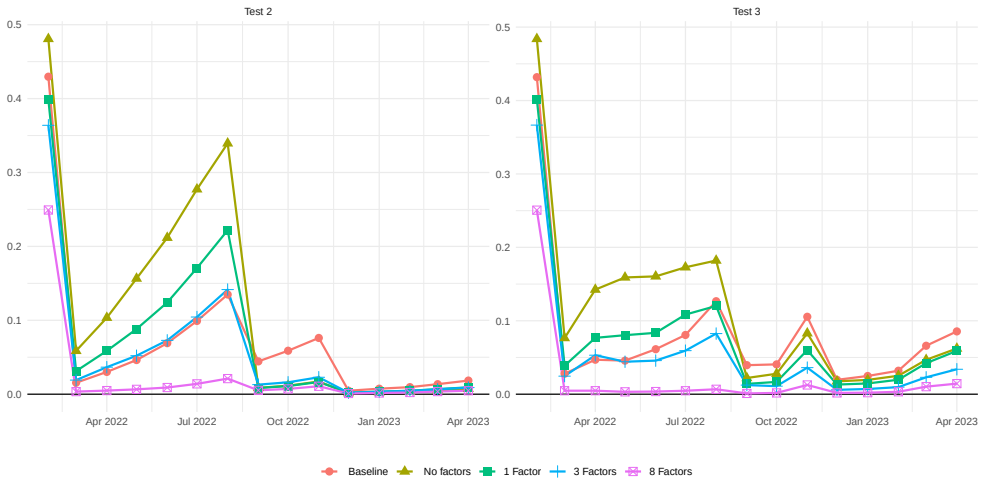


Figure A23: Hypothesis tests: Monetary policy shock, different number of factors

Figure shows the p-values of the rolling hypothesis tests that the impact of the monetary policy shock is unchanged in an expanding window starting in February 2022. Tests 2 and 3 test, respectively, whether the transmission of monetary policy has changed at some point and on average in the test window. Lines differ in the number of factors, which in the baseline is five.

B.4 Including speeches

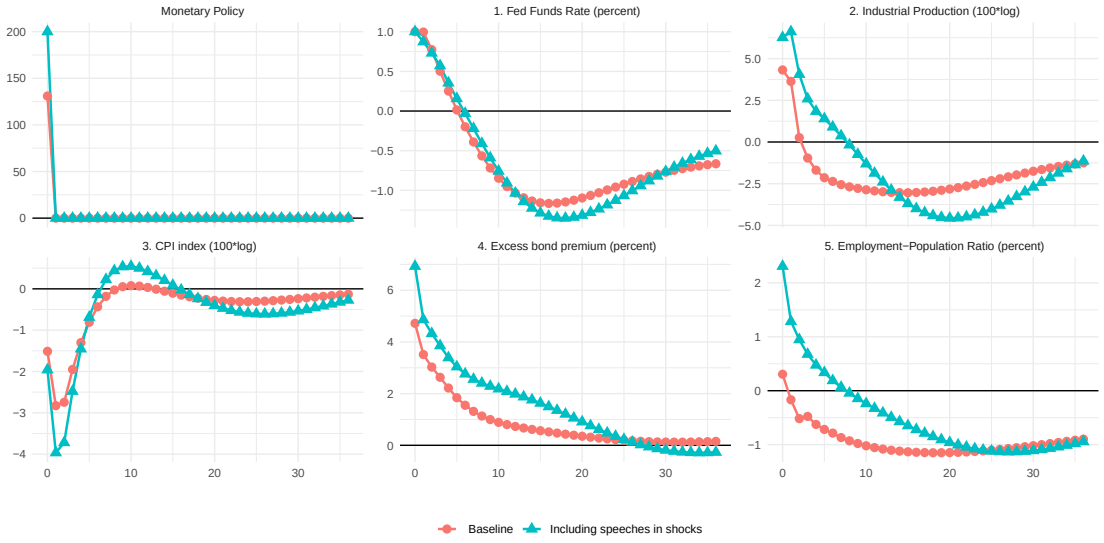


Figure A24: Impulse responses: Monetary policy shock, including speeches.

Figure shows the point estimates of the response of a factor-augmented VAR to a monetary policy shock, scaled to a 1 percentage point increase in the interest rate. Green lines use shocks including policymakers' speeches, while the baseline only includes shocks based on FOMC meetings. All specifications include five factors, responses of which are not shown here.

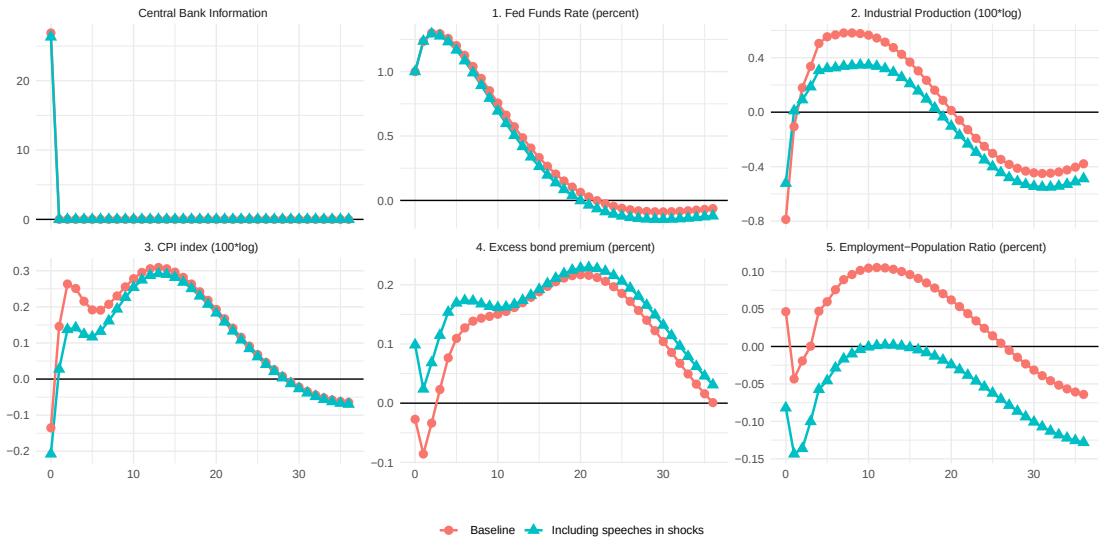


Figure A25: Impulse responses: Central Bank information shock, including speeches.

Figure shows the point estimates of the response of a factor-augmented VAR to a Central Bank information shock, scaled to a 1 percentage point increase in the interest rate. Green uses shocks including policymakers' speeches, while the baseline only includes shocks based on FOMC meetings. All specifications include five factors, responses of which are not shown here.

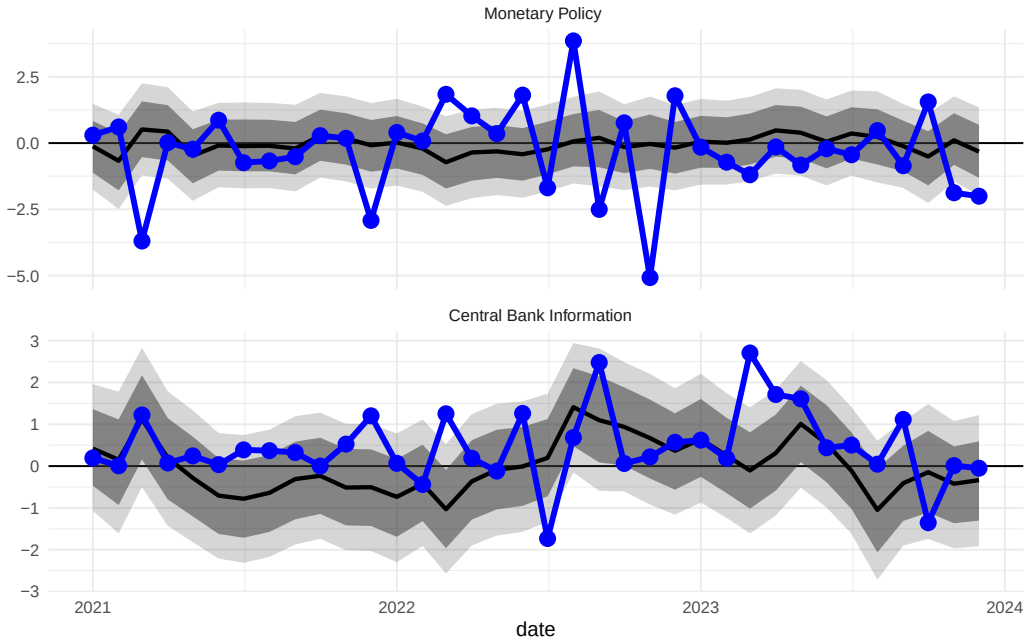


Figure A26: Inferred and realized shocks: 2021-2023, including speeches in shocks

Solid lines are point estimates of the inferred shocks under the null hypothesis that the data generating process remains unchanged. The shaded regions are the 68 and 90 percent confidence intervals from a bootstrap with $K = 1000$ replications. The blue line with large dots is the actual shock, computed from high frequency data. This version uses shocks including policymakers' speeches.

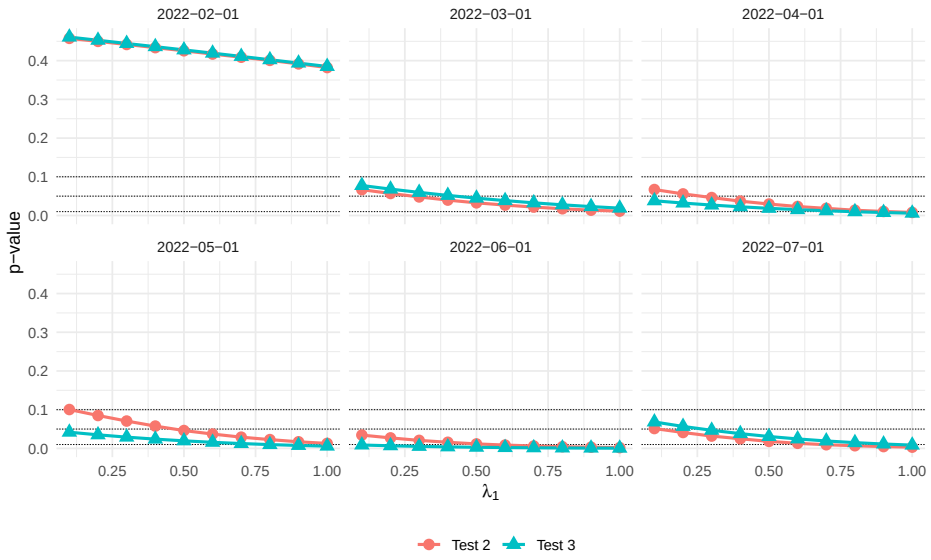


Figure A27: p-values for hypothesis tests as strength of monetary policy shock varies

Figure shows p-values for hypothesis tests described in Section 5.3, under a series of nulls, indexed by λ_1. Horizontal lines show 1, 5, and 10 percent confidence levels. The interpretation of λ_1 is that it represents a null where monetary policy is only λ_1 as effective as pre-COVID. Each panel computes the corresponding hypothesis tests on a window starting in February 2022 and ending in the titular month. The values for both lines when $\lambda_1 = 1$ thus match the corresponding columns in Table 6. This version uses shocks including policymakers' speeches.

Shock	$\mathcal{T}$ Start	$\mathcal{T}$ End	$\lvert\mathcal{T}\rvert$	Test 1	Test 2	Test 3
Monetary Policy	2022-02-01	2022-02-01	1	0.385	0.382	0.385
	2022-02-01	2022-03-01	2	0.385	0.011**	0.019**
	2022-02-01	2022-04-01	3	0.385	0.008***	0.006***
	2022-02-01	2022-05-01	4	0.385	0.013**	0.006***
	2022-02-01	2022-06-01	5	0.385	0.002***	0.001***
	2022-02-01	2022-07-01	6	0.933	0.003***	0.008***
Central Bank Information	2021-05-01	2021-05-01	1	0.281	0.281	0.281
	2021-05-01	2021-06-01	2	0.281	0.311	0.162
	2021-05-01	2021-07-01	3	0.281	0.192	0.061*
	2021-05-01	2021-08-01	4	0.281	0.162	0.029**
	2021-05-01	2021-09-01	5	0.281	0.182	0.022**
	2021-05-01	2021-10-01	6	0.399	0.234	0.026**
	2021-05-01	2021-11-01	7	0.399	0.185	0.013**
	2021-05-01	2021-12-01	8	0.399	0.070*	0.003***
	2021-05-01	2022-01-01	9	0.399	0.070*	0.002***
	2021-05-01	2022-02-01	10	0.507	0.091*	0.003***
	2021-05-01	2022-03-01	11	0.507	0.012**	0.000***
	2021-05-01	2022-04-01	12	0.507	0.016**	0.000***

Table A1: p-values for joint hypothesis tests, including speeches in shocks

Table shows p-values for the joint hypothesis that consecutive observations are drawn from the null. Each line considers three tests that the observed high frequency shocks are drawn from a distribution with a higher mean than would be expected based on pre-COVID transmission of monetary policy, over a period of consecutive months. $^{*}p < 0.1,^{**}p < 0.05,^{***}p < 0.01$. This version uses shocks including policymakers' speeches.

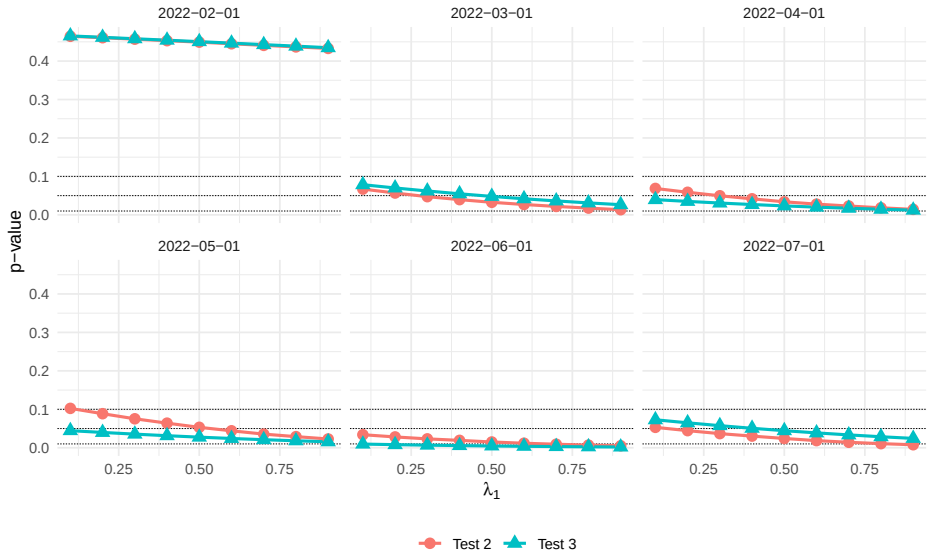


Figure A28: p-values for hypothesis tests as strength of monetary policy shock varies: Post-2021 residual covariance

Figure shows p-values for hypothesis tests described in Section 5.3, under a series of nulls, indexed by λ_1. Horizontal lines show 1, 5, and 10 percent confidence levels. The interpretation of λ_1 is that it represents a null where monetary policy is only λ_1 as effective as pre-COVID. Each panel computes the corresponding hypothesis tests on a window starting in February 2022 and ending in the titular month. The values for both lines when $\lambda_1 = 1$ thus match the corresponding columns in Table 6. This version uses the post-2021 residual covariance for $\hat{\Omega}_y$, and shocks including policymakers' speeches.

B.5 External instruments

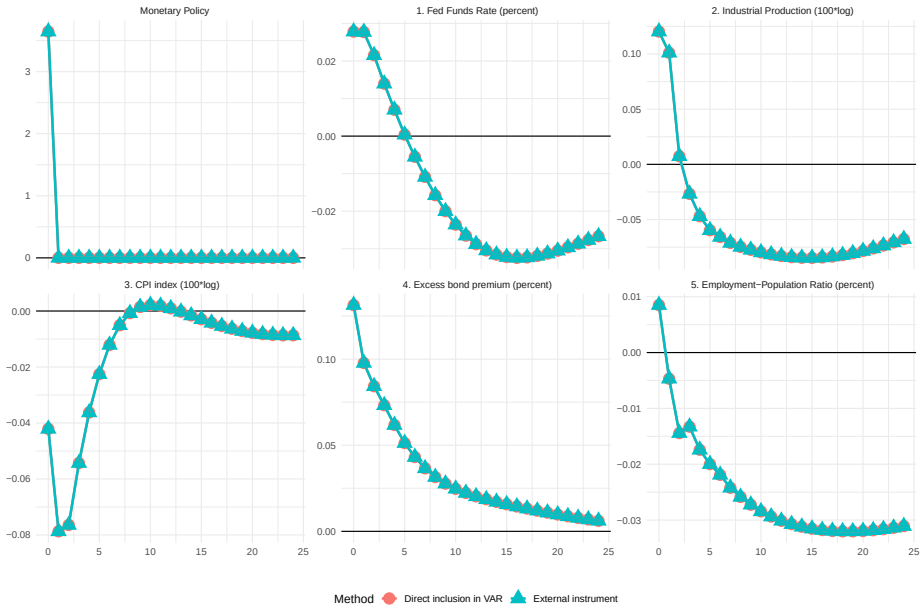


Figure A29: Direct inclusion versus external instruments: Monetary policy shock

Line labelled "External instrument" shows the point estimate of the response of the headline variables to a one-standard deviation shock computed using the high-frequency monetary policy series as an external instrument. The line labelled "Direct inclusion in VAR" are our baseline IRFs, rescaled to match the same initial impulse as for the external instrument. The five factors are included in the estimation but, for clarity, their responses are omitted.

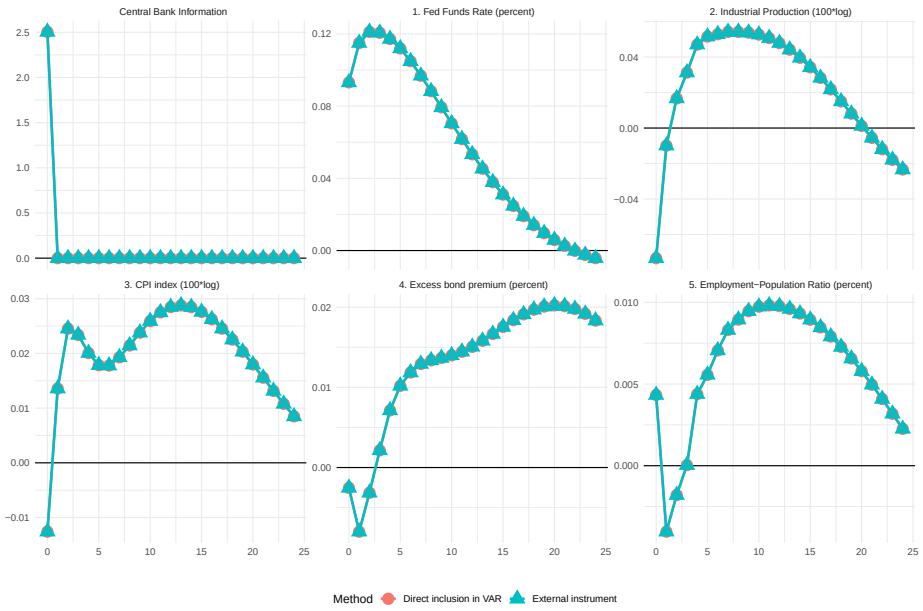


Figure A30: Direct inclusion versus external instruments: Central Bank information shock

Line labelled "External instrument" shows the point estimate of the response of the headline variables to a one-standard deviation shock computed using the high-frequency monetary policy series as an external instrument. The line labelled "Direct inclusion in VAR" are our baseline IRFs, rescaled to match the same initial impulse as for the external instrument. The five factors are included in the estimation but, for clarity, their responses are omitted.

B.6 Results from the structural model

Slope of	Baseline	Parameter changed	Test 2		Test 3	
			Acceptance region	Max p-value	Acceptance region	Max p-value
Phillips Curve	0.021	ξ_p, Price Calvo probability	(0.02, 0.035)	0.2	(0.019, 0.042)	0.19
Wage Phillips Curve	0.01	ξ_w, Wage Calvo probability	(0.0098, 0.017)	0.19	(0.0093, 0.018)	0.19
IS Curve	0.12	σ_c, Inter-termporal EoS	(0.12, 0.13)	0.15	(0.12, 0.14)	0.14

Table A2: Summary results: changing IS and Phillips curve slopes in a structural model, Smets and Wouters (2007) calibration.

Table reports the results of perturbing the Smets and Wouters (2007) model of the US economy by various parameters and using the corresponding impulse response functions in our tests. The model is calibrated based on Smets and Wouters (2007), see Appendix D.2.

C Deriving the hypothesis tests

Here we derive the hypothesis tests we apply in Section 5.3.

C.1 Set up

Recall that η_t is the difference between the observed high-frequency shocks and their filter-inferred equivalent, given by $\eta = v_t - \bar{\delta}_t$ where $\bar{\delta}_t = \int_\delta \delta f_t(\delta) d\delta$ and $(\sigma_t^i)^2 = \int_\delta (\delta - \bar{\delta}_t)^2 f_t(\delta) d\delta$ are the mean and variance of the filter-implied distribution for the ith structural shock.

We define the parameter $\mu_t = \mathbb{E}_t \eta_t$ as the mean of η_t and want to test jointly whether some or all of these parameters are positive on a sample $t = 1, \ldots, T$. Specifically, we want to test:

Test 1 $H_0 : \mu_t^i \leq 0$ for some $t \in \mathcal{T}$ vs. $H_1 : \mu_t^i > 0$ for all $t \in \mathcal{T}$

Test 2 $H_0 : \mu_t^i \leq 0$ for all $t \in \mathcal{T}$ vs. $H_1 : \mu_t^i > 0$ for some $t \in \mathcal{T}$

Test 3 $H_0 : \bar{\mu}^i \leq 0$ vs. $H_1 : \bar{\mu}^i > 0$, where $\bar{\mu}^i = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \mu_t^i$

It is convenient to fix i and work with the normalized shock $x_t = \eta_t^i / \sigma_t^i$, where σ_t^i is the root mean square error of the filter-inferred shock. As is common, we assume that this is known. Normality of the filter-inferred estimator implies that:

$$x_t \sim_{i,i,d} N(\mu_t, 1) \quad \forall\, t \in \mathcal{T}$$

Given this, we can deal with test 3 straightforwardly, since it becomes a standard univariate one-sided test of the sample mean. That is, we can construct the test statistic

$$\bar{x}_{\mathcal{T}} = \frac{1}{|T|} \sum_{t \in \mathcal{T}} x_t$$

Then

$$\mathbb{E}\bar{x}_{\mathcal{T}} = \bar{\mu}^i = \frac{1}{T} \sum_{t \in \mathcal{T}} \mu_t^i$$

where $T = |\mathcal{T}|$. And so

$$\bar{x}_{\mathcal{T}} \sim N(\bar{\mu}^i, 1/\sqrt{T})$$

Under the null for test 3, the mean of this distribution is zero, so we can compute one-sided p-values for the data from:

$$p_3(\bar{x}_{\mathcal{T}}) = 1 - \Phi(\bar{x}_{\mathcal{T}}/\sqrt{T})$$

However, for tests 1 and 2, things are less straightforward. We derive likelihood ratio tests more formally in the following sections.

C.2 Rejection Regions

Tests 1 and 2 both consider the tests of the form $H_0 : \theta \in \Theta_0$ versus $H_1 : \theta \in \Theta_0^c$, where Θ_0 is a restricted subset of the parameter space Θ and Θ_0^c is the complement of Θ_0.

The likelihood ratio test statistic is then:

$$\lambda(\boldsymbol{x}) = \frac{\sup_{\Theta_0} L(\theta|\boldsymbol{x})}{\sup_{\Theta} L(\theta|\boldsymbol{x})} \tag{9}$$

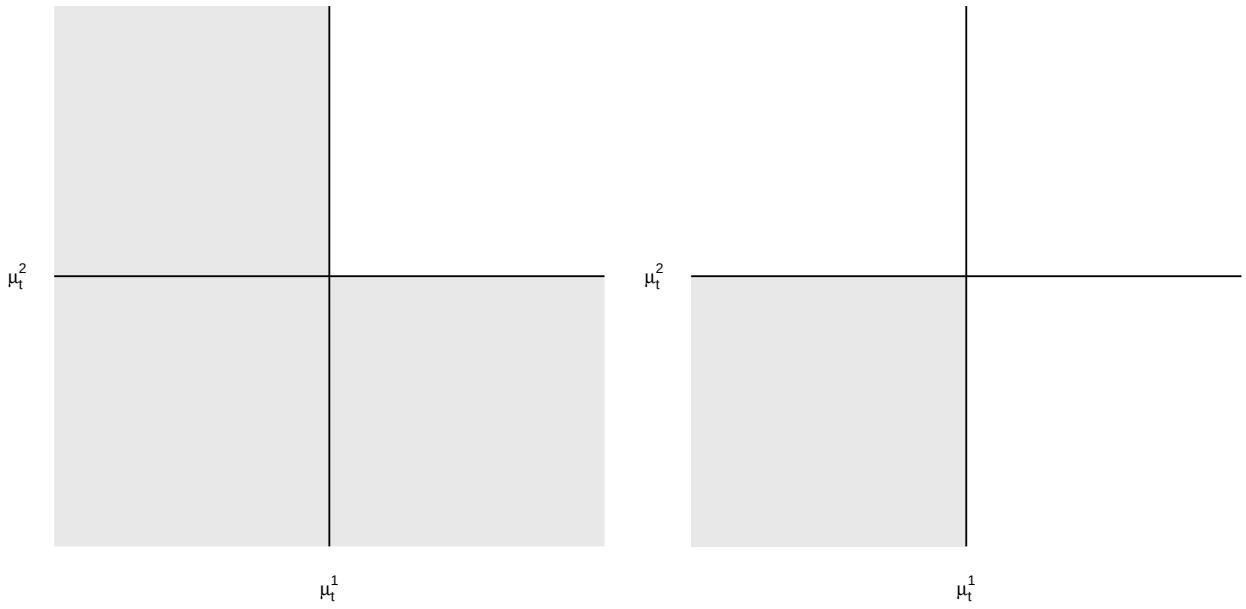
where $\boldsymbol{x} = (x_1, \ldots, x_T)$ is the data and $L(\cdot)$ is the likelihood function.[22]

In our test 1, the parameter restriction is that $\mu_t^i \leq 0$ for some $t \in \mathcal{T}$. This is a fairly loose restriction, and to reject corresponds to the idea that one is certain that monetary policy is less effective than usual in all the periods in $\mathcal{T}$. We illustrate this for the simple case where $\mathcal{T} = \{1, 2\}$ in Figure A31a. In contrast, the parameter restriction for test 2 is much tighter, requiring that $\mu_t^i \leq 0$ for all $t \in \mathcal{T}$. This corresponds to the idea that monetary policy was less effective than usual at least some of the time during $\mathcal{T}$. In the two-period case, this is shown by the smaller shaded region in Figure A31b.

[22] See chapter 8 of Casella and Berger (2024) for a textbook treatment of this topic, which we follow closely here.

Formally, we define the parameter restriction sets for the two tests:

$$\Theta_0^1 = \{\boldsymbol{\mu} \in \mathbb{R}^T \mid \mu_t \leq 0 \ \ \forall\, t \in \mathcal{T}\}$$
$$\Theta_0^2 = \{\boldsymbol{\mu} \in \mathbb{R}^T \mid \mu_t \leq 0 \text{ for some } t \in \mathcal{T}\}$$



(a) Test 1 (b) Test 2

Figure A31: Parameter restrictions for tests 1 and 2

Note: Shaded regions show the restricted parameter under the null, Θ_c, for tests 1 and 2 for the case where $\mathcal{T} = \{1, 2\}$.

The likelihood ratio test rejects H_0 in favor of H_1 when $\lambda(\boldsymbol{x})$ falls below some critical value. For any given critical value, the set of possible observations of the data which would cause the test to reject the null is known as the rejection region. Because the data are drawn from independent normal distributions with unit variance, the likelihood ratio is simply a function of the minimum Euclidian distance from the observed data to the nearest corresponding point inside Θ_c.

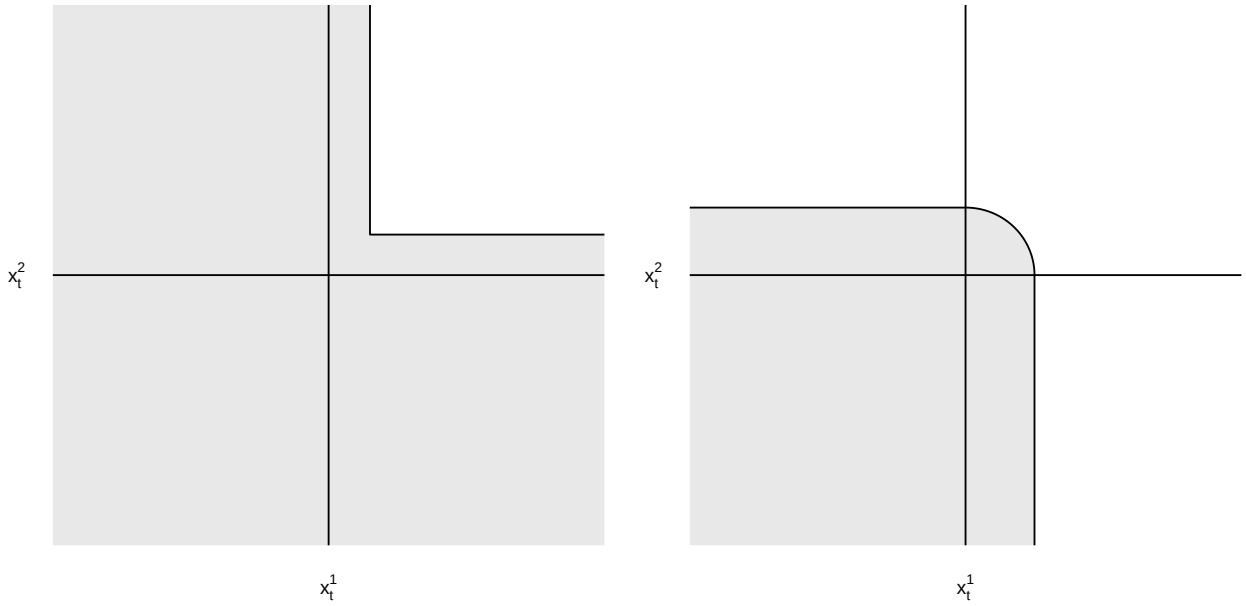
For each test $j = 1, 2$, we can write down the general formula for a rejection regions with a given point $\boldsymbol{x}$ on the boundary as:

$$R_j(\boldsymbol{x}) = \{\boldsymbol{X}' \in \mathbb{R}^t \mid d_j(\boldsymbol{x}') \geq d_j(\boldsymbol{x})\}$$

where $d_j(\boldsymbol{x})$ is the Euclidian distance between $\boldsymbol{x}$ and the nearest point inside the constraint set:

$$d_j(\boldsymbol{x}) = \inf_{\boldsymbol{\mu} \in \Theta_0^j} ||\boldsymbol{x} - \boldsymbol{\mu}||$$

This produces rejection regions which look like the parameter restriction sets plus a buffer (see Figure A32 for the two-period example once more). This buffer is larger for smaller critical values of the likelihood ratio. More intuitively, the critical value for the likelihood ratio test sets the bar for how much the data have to disagree with the null sufficiently in order to reject it. Of course, likelihood ratio critical values are not easy to interpret. In the next section, we derive the more familiar p-values for these tests.



(a) Test 1 (b) Test 2

Figure A32: Rejection regions for tests 1 and 2

Note: Unshaded regions show rejection regions for tests 1 and 2 for the case where $\mathcal{T} = \{1, 2\}$.

C.3 Computing p-values

The p-value is a function of the data, $p(\boldsymbol{x})$. This is usually described as the probability of rejecting the observed data under the null. However, with multivariate set nulls (as we have here) this definition is insufficient. Instead, a valid p-value is the maximum probability of

rejection over the null set, akin to a worst-case scenario (again see Casella and Berger (2024) for details).

$$p_j(\boldsymbol{x}) = \sup_{\boldsymbol{\mu} \in \Theta_0^j} \mathbb{P}(\boldsymbol{x}' \in R_j(\boldsymbol{x}))$$

For each test, this means we can compute the p-value in two steps. First, we calculate $\boldsymbol{\mu}_j^* \in \Theta_0^j$, the parameter value which maximizes this probability over the null set. Second, we then integrate over draws of the data conditional on $\boldsymbol{\mu}^*$.

For test 1, we can compute a closed-form solution for this test:

$$p_1(\boldsymbol{x}) = 1 - \Phi(\min_{t \in \mathcal{T}} x_t)$$

However, for test 2, we can only complete the first step analytically, since $\boldsymbol{\mu}_j^* = \mathbf{0}$, the origin. To compute the p-values, we integrate numerically, drawing 100,000 points from the mean-zero unit-variance multivariate normal and computing the fraction of observations in $R_2(\boldsymbol{x})$.

D Structural Models

In this section we describe two different structural models and discuss how they can be mapped into our empirical exercises. The first is intentionally simple, and we use it only to illustrate explicitly the point that weaker transmission of monetary policy shocks can have different impacts on different variables depending on how that weaker transmission is implemented (i.e. which structural parameter is changed to deliver it). In the second, we outline a more complicated structural economic model which we actually apply to the data.

D.1 A simple New Keynesian model

Consider the standard linear New Keynesian Model where inflation π_t, the output gap y_t, the nominal interest rate i_t, and the natural interest rate r_t^n are related by:

$$\begin{aligned}
\pi_t &= \beta \mathbb{E}_t \pi_{t+1} + \kappa y_t && \text{Phillips Curve} \\
y_t &= -\frac{1}{\sigma}\left(i_t - \mathbb{E}_t \pi_{t+1} - r_t^n\right) + \mathbb{E}_t y_{t+1} && \text{Intertemporal IS curve} \\
i_t &= \rho + \phi_\pi \pi_t + \phi_y y_t + v_t && \text{Taylor Rule} \\
r_t^n &= \rho + \sigma \psi_{ya} \rho_a a_t && \text{Natural rate of interest}
\end{aligned}$$

The model is driven by exogenous shocks to monetary policy v_t and technology a_t. These are assumed to follow first-order autoregressive processes:

$$v_t = \rho_v v_{t-1} + \epsilon_t^v \qquad\qquad a_t = \rho_a a_{t-1} + \epsilon_t^a$$

Uniqueness of the saddle path solution is guaranteed if $\kappa(\phi_\pi - 1) + (1-\beta)\phi_y > 0$. For further details and derivations, the reader is directed to Chapter 3 of Galí (2015), which also solves explicitly for the contemporaneous response to a monetary policy shock. That is:

$$D_m(\theta) = \begin{bmatrix} \frac{\partial y_t}{\partial v_t} \\ \frac{\partial \pi_t}{\partial v_t} \end{bmatrix} = \begin{bmatrix} -(1-\beta\rho_v)\Lambda_v \\ -\kappa\Lambda_v \end{bmatrix}$$

Where

$$\Lambda_v = \left((1-\beta\rho_v)\left(\sigma(1-\rho_v)+\phi_y\right)+\kappa\left(\phi_\pi-\rho_v\right)\right)^{-1}$$

These derivatives are the initial points of the impulse response to a monetary policy shock. Under the parameter restrictions required for uniqueness, both entries are negative; a positive monetary shock reduces both inflation and the output gap on impact.

We consider two ways in which monetary transmission may be weakened: either a flatter Phillip Curve (κ falls), or a flatter intertemporal IS curve (σ rises). The relevant derivatives are:

$$\frac{\partial D_m(\theta)}{\partial \kappa} = \begin{bmatrix} (1-\beta\rho_v)\Lambda_v^2\phi_\pi \\ \Lambda_v(1-\kappa\phi_\pi\Lambda_v) \end{bmatrix} \qquad \frac{\partial D_m(\theta)}{\partial \sigma} = \begin{bmatrix} (1-\beta\rho_v)^2\Lambda_v^2(1-\rho_v) \\ \kappa(1-\beta\rho_v)\Lambda_v^2(1-\rho_v) \end{bmatrix}$$

The derivatives with respect to σ are both positive, so the initial impulse response is attenuated when the IS curve flattens (σ rises). Yet the impact on inflation and interest rates differs by a ratio of $\kappa(1-\beta\rho_v)$. This difference arises because the Phillips curve induces a positive correlation between output and inflation. Thus, the simple assumption that weaker monetary transmission affects all variables equally is clearly violated if implemented by a flattening of the IS curve. Likewise, when the slope of the Phillips flattens (i.e. κ declines), the marginal effect on the impulse responses of inflation and the output gap differ. However, this case is even more extreme in that if ρ_v and $\kappa\phi_\pi$ are sufficiently large, the signs of the derivatives may differ. That is, a flatter Phillips curve can lead to an attenuation of the inflation response to a monetary tightening but an amplification of the output decline.

In sum, the preceding example shows how even in a simple model the impact of weaker monetary transmission can have different effects on different variables. This motivates the more general case discussed in Section 6.4.

D.2 A more complicated model

We use the canonical medium-scale DSGE model as in Smets and Wouters (2007). The model doesn't allow for an analytical solution. We can write the solution of a linearized DSGE model in the following state-space representation:

$$X_t = B(\theta)X_{t-1} + \Gamma(\theta)\eta_t \tag{10}$$

$$Y_t = A(\theta)X_t \tag{11}$$

Here X_t are predetermined, or state, variables, Y_t are observable variables, including non-predetermined, or jump, variables, A, B, and Γ are coefficient matrices, that depend on deep parameters θ, η_t is a vector of shocks, with $E(\eta_t) = 0$ and $E(\eta_t\eta_t') = I$. Under certain conditions, the state-space representation of the DSGE model has an infinite order VAR process representation:

$$Y_t = \sum_{j=1}^{\infty} \Delta_j Y_{t-j} + A\Gamma\eta_t \tag{12}$$

$$\Delta_j = ABM^{j-1}\Gamma(A\Gamma)^{-1} \tag{13}$$

$$M = [I_n - \Gamma(A\Gamma)^{-1}A]B \tag{14}$$

Equations (12)-(14) show how a solution to a DSGE model in state-space representation relates to a VAR model as we use in the main body of the text.

To estimate the DSGE model, we use an Impulse Response Matching approach as in Christiano et al. (2005). For this approach, we distinguish between two sets of parameters $\theta = \{\theta^1, \theta^2\}$. We set θ^1 equal to values taken from the literature. θ^2 are set to match the impulse response for a monetary policy shock from our VAR as closely as possible. Formally, our estimator of θ^2 is

$$J = \min_{\theta^2} [\hat{\Psi} - \Psi(\theta^2)]' V^{-1} [\hat{\Psi} - \Psi(\theta^2)] \tag{15}$$

where $\Psi(\theta^2)$ is the mapping from θ^2 to the structural model's impulse response function, $\hat{\Psi}$ is the impulse response function estimated using the VAR, and V is a weighting matrix.

The set of parameters that enter the impulse response matching algorithm are $\theta^2 \in \{\sigma_c, \xi_p, \xi_w, \phi_p, \varphi, \lambda, \psi, \sigma_L, \rho, \iota_p, \iota_w\}$. Here, σ_c is the inverse inter-temporal elasticity of substitution, ξ_p is the degree of price stickiness, ξ_w is the degree of wage stickiness, ϕ_p is one plus the share of fixed costs in production, φ is the elasticity of the capital adjustment cost function, λ is the habit parameter in consumption, ψ is a parameter related to the capital utilisation adjustment cost function, σ_L is the elasticity of labor supply with respect to the real wage, ρ captures the degree of interest rate smoothing in the monetary policy reaction function, ι_p is the degree of price indexation, and ι_w is the degree of wage indexation. This set of parameters is mostly the same as in Christiano et al. (2005), except that we add ι_p and ι_w, as these parameters are not featured in the model of Christiano et al. (2005), and we add ρ as it better helps to match the response of the nominal interest rate.

The remaining parameters, θ^1, are set to the posterior means in Smets and Wouters (2007), Table 1a. Exceptions are discount factor β and depreciation rate δ, which are set to imply a steady state annualized nominal interest rate of 3 percent and annualized depreciation rate of 10 percent, respectively. Furthermore, since our model is at monthly frequency, some additional parameters (γ, ρ_R) are transformed accordingly from the value in Smets and Wouters (2007).[23] Table A3 shows the structural parameters both for the version of the model where we use the impulse-response matching to estimate selected parameters (θ^2), and the parametrization of the original Smets and Wouters (2007) model.[24]

[23] The Smets and Wouters (2007) model is estimated at quarterly frequency, while the data used to generate our VAR impulse responses is at monthly frequency. Therefore, for the original Smets and Wouters (2007) parametrization, we change parameters β, δ, γ, and ρ_R accordingly.

[24] We use Dynare to solve the DSGE model and perform the impulse response matching.

Parameter	Symbol	Value est.	Value S&W
Elasticity of the capital adjustment cost function	φ	3.97	5.74
Coefficient of relative risk aversion	σ_c	0.28	1.38
Habit persistence parameter	λ	0.34	0.71
Degree of wage stickiness	ξ_w	0.13	0.70
Elasticity of labor supply with respect to the real wage	σ_L	2.78	1.83
Degree of price stickiness	ξ_p	0.54	0.66
Degree of wage indexation	ι_w	0.78	0.58
Degree of price indexation	ι_p	0.38	0.24
Parameter related to capital utilization adjustment cost function	ψ	0.36	0.54
1 + share of fixed cost in production	ϕ_p	1.73	1.60
Inflation parameter in monetary policy (MP) reaction function	r_π	2.04	2.04
Interest rate smoothing parameter in MP reaction function	ρ	0.64	0.81
Output gap parameter in MP reaction function	r_y	0.08	0.08
Change in output gap parameter in MP reaction function	$r_{\Delta y}$	0.22	0.22
Discount factor	β	0.9975	0.9975
Steady state growth rate	γ	0.14	0.14
Output elasticity of capital	α	0.19	0.19
Depreciation rate	δ	0.0083	0.0083
Kimball goods market aggregator	ϵ_p	10	10
Kimball labor market aggregator	ϵ_w	10	10
1 + steady state labor market mark-up	ϕ_w	1.5	1.5
Persistence of monetary shock	ρ_R	0.53	0.53

Table A3: Parametrization of structural model

Table reports parametrization of Smets and Wouters (2007) model of the US economy. Column "Value est." refers to version where selected parameters are estimated using impulse-response matching, "Value S&M" to parametrization in Smets and Wouters (2007), Table 1a.