Université de Neuchâtel
Institut de recherches économiques

IRENE Working Paper 24-03

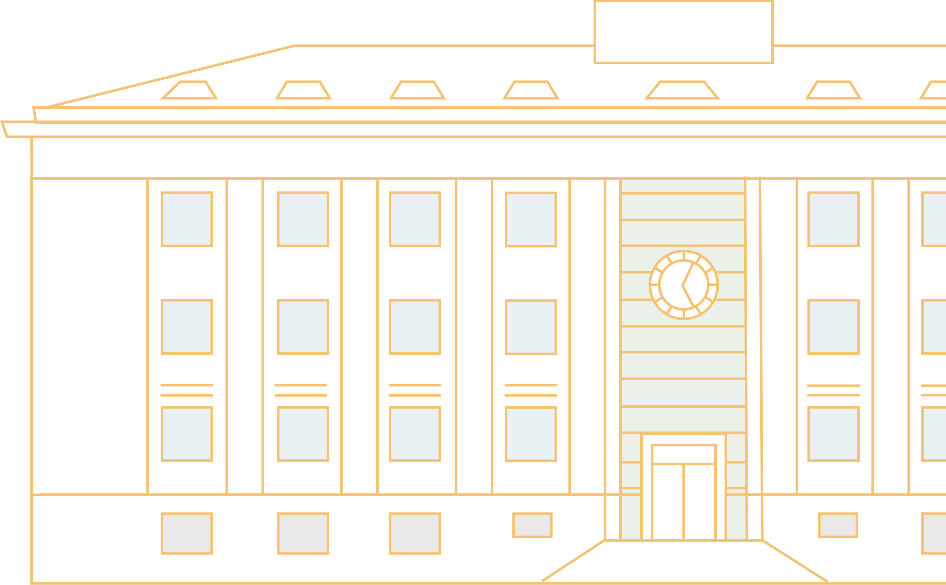

Measuring monetary policy shocks

Marc Burri and Daniel Kaufmann

Measuring monetary policy shocks[*]

Marc Burri[†] Daniel Kaufmann[‡][§]

2 April 2026

Abstract: We propose a two-step approach to measure monetary policy shocks based on daily financial market data. First, we estimate the causal impact of a monetary policy shock on financial market variables using standard instrumental variables techniques (high-frequency and heteroskedasticity-based identification). Second, we exploit the cross-sectional variation of the causal impact to predict the underlying unobserved monetary policy shocks based on the Kalman filter. The two-step approach delivers a more accurate measure of monetary policy shocks. As a consequence, various anomalies documented in the literature on measuring monetary policy shocks are alleviated or resolved.

JEL classification: C3, E3, E4, E5

Keywords: Monetary policy shocks, high-frequency identification, identification through heteroskedasticity, instrumental variables, weak-instrument problem, Kalman filter

---

[*] We thank Gianluca Benigno, Kenza Benhima, Rui Esteves, Aurélien Eyquem, Jean-Marie Grether, Margaret Jacobson, Kevin Kotzé, Leif Anders Thorsrud, Jean-Paul Renne, Eric Swanson, Cédric Tille, Mark Watson, Linyan Zhu, as well as participants at the SSES Annual Congress, the Workshop on Applied Macroeconomics and Monetary Policy in St. Gallen, the University of Fribourg Research Seminar, the University of Lausanne research seminar, and the Geneva Graduate Institute research seminar, for helpful discussions and comments. Daniel Kaufmann gratefully acknowledges financial support from the Swiss National Bank. A previous version of this paper circulated under the title "Multi-dimensional monetary policy shocks based on heteroscedasticity".

[†] University of Neuchâtel, Institute of Economic Research, Rue A.-L. Breguet 2, CH-2000 Neuchâtel, marc.burri@unine.ch

[‡] University of Neuchâtel, Institute of Economic Research, Rue A.-L. Breguet 2, CH-2000 Neuchâtel, daniel.kaufmann@unine.ch

[§] KOF Swiss Economic Institute, ETH Zurich

1 Introduction

Identifying the causal effects of monetary policy is a challenge. Most decisions taken by central banks are endogenous to the state of the economy. Researchers exploit quasi-random variation in financial market data to circumvent this problem. High-frequency identification schemes construct monetary policy surprises using changes in financial market variables in a narrow window around central bank announcements (see, e.g., Kuttner, 2001, Gürkaynak et al., 2005, Altavilla et al., 2019, Swanson, 2021). These surprises can then be used as external instruments to identify causal effects on low-frequency macroeconomic variables (see, e.g., Gertler and Karadi, 2015, Stock and Watson, 2018). Heteroskedasticity-based identification schemes exploit the change in the variance-covariance of financial market variables on days with monetary policy announcements to measure their effect on financial market variables (see, e.g., Rigobon, 2003, Rigobon and Sack, 2004, Bu et al., 2021).

The literature has documented various anomalies associated with those identification strategies: High-frequency surprises are predictable based on information known before the announcement (see Bauer and Swanson, 2022, 2023, Miranda-Agrippino and Ricco, 2021, Zhu, 2023), show unexpected co-movements with stock prices (see Nakamura and Steinsson, 2018, Miranda-Agrippino and Ricco, 2021), and suffer from weak-instrument problems (see Ramey, 2016, Bauer and Swanson, 2022). In addition, Brennan et al. (2024) show that the correlation between various versions of high-frequency surprises, but also with shocks based on heteroskedasticity by Bu et al. (2021), is low. It is an open question whether different approaches identify different shocks or whether the issue lies with the precision of the shock measurements.

We approach this question by treating the monetary policy shocks as unobserved objects that have to be predicted from the data. It is widely recognized that high-frequency surprises are proxies of underlying unobserved shocks (see, e.g., Stock and Watson, 2018). In heteroskedasticity-based identification schemes, the underlying shocks are completely unobserved because we exploit the change in the variance-covariance matrix to identify the impact matrix (see Bu et al., 2021). Some of the anomalies may therefore stem from the fact that existing approaches do not accurately measure the underlying monetary policy shocks.

This paper proposes a two-step approach that allows us to predict the underlying unobserved shocks from the cross-sectional impact of monetary policy shocks on daily financial variables. In a first step, we identify the causal impact of monetary policy shocks on a variety of daily financial market variables using standard high-frequency and heteroskedasticity-based instruments (see, e.g., Nakamura and Steinsson, 2018, Rigobon and Sack, 2004).

In a second step, we derive the optimal prediction of the unobserved shocks, which minimizes the mean squared error (MSE). A simulation exercise shows that, under realistic assumptions on the data generating process and sample size, the predictions are more highly correlated with the true underlying shocks than a noisy proxy. Intuitively, the IV-estimates of the daily impact on financial market variables are accurate due to the large number of daily observations available. In addition, the predictions of the underlying shocks are accurate because of the relatively large cross-section of financial market variables. Therefore, the prediction error is less important than the noise that is filtered out of the proxy.

We then apply the approach to US data from 1988–2019. First, we examine whether heteroskedasticity and high-frequency identification schemes differ. We find that the estimated impact on various financial variables is qualitatively similar. Therefore, the predicted shocks are highly correlated (0.82 and higher). This suggests that the low correlation reported by Brennan et al. (2024) can be traced in part back to measurement error in the high-frequency surprises. Second, we examine the weak-instrument problem documented by Ramey (2016). We show that the predicted shocks yield stronger instruments in monthly models for macroeconomic data than when using the high-frequency surprises directly. Again, this is consistent with the view that high-frequency surprises are contaminated by measurement error. The strongest instrument we find is for heteroskedasticity-based monetary policy shocks. Third, we investigate whether two-step predictions are correlated with the information available at the time of the announcement, a phenomenon documented for high-frequency surprises by Bauer and Swanson (2023, 2022). The explanatory power of the predictors by Bauer and Swanson (2022) is substantially reduced when using the two-step approach. The reason is that the prediction incorporates control variables. Therefore, an ex-post orthogonalization of the monetary policy shocks is not necessary. Again, the heteroskedasticity-based two-step predictions perform slightly better. Third, we investigate the anomaly that high-frequency surprises often move in the same direction as stock prices. This phenomenon can potentially be explained by the information the central bank releases when a policy decision is announced (see Nakamura and Steinsson, 2018, Jarociński and Karadi, 2020). However, we show that the share of observations with an unexpected positive co-movement is 10 percentage points lower for the two-step predictions than the high-frequency surprises. This suggests that part of the co-movement may be because high-frequency surprises are contaminated by measurement error. Finally, we examine whether the two-step predictions are a useful guide for policy makers in real-time. We perform a pseudo-out-of-sample forecasting exercise and find that the real-time predictions of the shocks are highly correlated with the final, that is end-of-sample, predictions.

We relate to existing papers that aim to measure monetary policy shocks. A sufficient condition to recover the underlying shocks in a vector autoregression (VAR) is that the model is invertible (see Fernández-Villaverde et al., 2007, Stock and Watson, 2018). However, this is a strong assumption which unlikely holds in practice. We propose a two-step approach that delivers the optimal prediction in a minimum-MSE sense if this assumption does not hold. Bu et al. (2021) propose to extract monetary policy shocks regressing interest rate changes across the term structure on the impact responses in the spirit of Fama and MacBeth (1973). This approach requires assumptions that may be violated in empirical applications. Specifically, it requires an orthogonality condition between the impact response to monetary policy shocks and the impact response to other shocks. If this condition is violated, the monetary policy shock series will be contaminated by other shocks occurring on policy event days. Our approach delivers the optimal prediction even if this assumption is not fulfilled. Finally, Bauer and Swanson (2023, 2022) propose to orthogonalize high-frequency surprises with respect to the information available at the time of the announcement. However, this does not account for unpredictable intraday noise in the surprises.[1] The two-step approach incorporates control variables when estimating the impact effect and when predicting the unobserved underlying shocks. In addition, it removes unpredictable intraday noise.

The remainder of the paper is structured as follows. Section 2 presents the two-step approach to measuring monetary policy shocks. Section 3 performs a simulation exercise to examine the conditions under which the two-step predictions are more accurate than a noisy proxy of the unobserved shocks. Section 4 examines the performance of the two-step approach in three empirical applications using US data. The last section concludes.

2 Measuring monetary policy shocks

Suppose that we observe N financial market variables (y_t), for which the data-generating process reads:

$$\begin{aligned} y_t &= \alpha + \Psi \varepsilon_t + \Gamma v_t + \Phi(L) x_{t-1} \quad \text{for } t \in P \\ y_t &= \alpha + \Gamma v_t + \Phi(L) x_{t-1} \quad \text{for } t \in C \ , \end{aligned} \tag{1}$$

---

[1] In addition, including irrelevant regressors in the orthogonalization step may introduce noise in small samples. However, the same problem arises in the two-step approach; We defer the question of variable selection to future research.

where ε_t is a vector of E structural shocks on policy event days (P), and v_t is a vector of R other shocks on policy event as well as control days (P and C).[2] We assume that all shocks are serially and mutually uncorrelated. Furthermore, Γ and Ψ denote impact matrices of dimensions $N \times R$ and $N \times E$, respectively. Finally, α is a vector of constant terms and x_{t-1} is a vector of pre-determined control variables, which may include lags of y_t, and $\Phi(L)$ is a conformable lag polynomial.

The first step consists of estimating one or more columns of Ψ up to a scale.[3] Following the seminal work of Kuttner (2001), researchers use intraday tick data to measure changes in financial market variables in a narrow window around central bank announcements (see, e.g., Gürkaynak et al., 2005, Nakamura and Steinsson, 2018, Jarociński and Karadi, 2020, Altavilla et al., 2019, Swanson, 2021, Bauer and Swanson, 2022). These high-frequency surprises can be used to approximate one or more structural shocks ε_t. If these surprises capture only part of the variation of ε_t, or are contaminated by unpredictable noise, one or more columns of Ψ can still be estimated using IV up to a scale (see, e.g., Stock and Watson, 2018, and Appendix A).

Another strand of the literature exploits the fact that the variance-covariance of y_t changes on policy event days compared to control days. Following Rigobon (2003), researchers exploit financial market data to estimate one column of Ψ using IV (see Rigobon and Sack, 2004, Lewis, 2022, and Appendix A) or multiple columns using GMM (Rigobon, 2003, Rigobon and Sack,

---

[2] We assume that policy event days occur only on policy event days for ease of exposition. Depending on the identification strategy, weaker assumptions suffice. In heteroskedasticity-based identification schemes we only need that the variance of the shocks is larger on policy event days (see Rigobon, 2003). In high-frequency identification schemes, P includes the days where high-frequency surprises are measured. However, we can still estimate Ψ if there are policy event shocks on days where we do not measure the high-frequency surprises, because those observations are not used in the IV estimator (see Appendix A).

[3] See Appendix A for a technical discussion.

2004, Lewis, 2022).[4]

The two approaches have advantages and disadvantages. High-frequency identification requires intraday tick data and knowledge of the exact timing of the announcement. Although a narrower event window makes the identification assumption that monetary policy shocks dominate over other shocks occurring on the same day, more credible, researchers may miss relevant variation if monetary policy shocks affect financial market variables with some delay.[5] While heteroskedasticity-based identification does not necessarily require intraday tick data, we need to assume that the variance-covariance matrix changes on event days. This assumption may not be fulfilled, for example, if other events occur systematically on control days but not on policy event days. In addition, heteroskedasticity-based identification schemes are sensitive to non-linearities in the data-generating process (see Kolesár and Plagborg-Møller, 2025).

The common feature of both approaches is that even if we can estimate Ψ, the underlying shocks ε_t are unobserved. In a second step, we therefore obtain the linear minimum-MSE prediction of the unobserved monetary policy shocks based on the Kalman filter.[6] In doing so, we treat the monetary policy shocks as unobserved or imperfectly observed objects that must be predicted based on available information.

Let $\varepsilon_{t|T}$ denote the prediction of the unobserved shocks based on all available information and

---

[4] Canetg and Kaufmann (2022) and Burri and Kaufmann (2025) additionally impose recursive zero-restrictions on a heteroskedasticity-based identification scheme to disentangle target and forward guidance shocks.

[5] Casini and McCloskey (2025) provide precise conditions under which high-frequency identification schemes identify the causal effects in a narrow time window.

[6] We thank Mark Watson for guiding us in this direction.

Σ the variance-covariance matrix of the residuals $u_t = y_t - \alpha - \Phi(L)x_{t-1}$.

Proposition 1. *For Model (1), the minimum-MSE prediction of ε_t for $t \in P$ is*

$$\varepsilon_{t|T} = \Sigma_\varepsilon \Psi' \Sigma^{-1} u_t$$

with MSE

$$MSE_{t|T} = \Sigma_\varepsilon - \Sigma_\varepsilon \Psi' \Sigma^{-1} \Psi \Sigma_\varepsilon$$

Proof. See Appendix B. □

Four comments are in order. First, we can usually estimate each column of Ψ only up to a scale. However, this only affects the scale of the shock prediction (see Appendix B). Second, we can estimate $\Sigma = \Psi\Sigma_\varepsilon\Psi' + \Gamma\Sigma_v\Gamma'$ as the variance-covariance matrix of u_t on policy event days ($t \in P$). Third, although we do not directly observe Σ_ε, it affects only the scale of the shocks. Therefore, we can apply an arbitrary normalization. Fourth, standard Kalman-filter results imply that the prediction is optimal if the shocks are normally distributed. However, even if they are not, we still obtain the best linear prediction in a MSE-sense.

2.1 Accuracy relative to high-frequency surprises

Every prediction is subject to prediction error. Whether the two-step prediction is more accurate than high-frequency surprises based on intraday tick data is therefore an empirical question. To see this, assume that we have only one structural shock ($E = 1$), one other shock ($R = 1$), and one variable to estimate the impact vector ($N = 1$). In what follows, lower-case letters denote the impact scalars and variances.

Suppose that we obtain a proxy of this shock: $Z_{1t} = \varepsilon_{1t} + \eta_{1t}$ for $t \in P$, where η_{1t} is an unpredictable noise term with variance $\sigma^2_{1\eta}$.[7] The precision of the proxy depends on the variance of the noise term, because the MSE corresponds to the variance of the prediction error: $\mathbb{V}(\varepsilon_{1t} - Z_{1t}) = \sigma^2_{1\eta}$.

How accurate is the two-step prediction? Assume that we know the impact scalar ψ from the first step. In a second step, the Kalman filter prediction reads:

$$\varepsilon_{1t|T} = \frac{\psi\sigma^2_\varepsilon}{\psi^2\sigma^2_\varepsilon + \gamma^2\sigma^2_v} u_t$$

with

$$MSE_{t|T} = \sigma^2_\varepsilon\left(1 - \frac{\psi^2\sigma^2_\varepsilon}{\psi^2\sigma^2_\varepsilon + \gamma^2\sigma^2_v}\right) \; .$$

The two-step prediction will become more accurate, that is, the MSE will become smaller if $\psi^2\sigma^2_\varepsilon$ increases relative to $\gamma^2\sigma^2_v$. This implies that the accuracy of the prediction depends on how strongly the financial market variables are driven by policy event shocks relative to other shocks on policy event days. In practice, the MSE will also be affected by estimation error and therefore depends on the number of observations and events in the sample. In addition, cross-sectional variation in the response of multiple variables may provide additional useful information to predict the shocks. Therefore, in practice, increasing N can also increase the

---

[7] Therefore, it may include classical measurement error. The classical assumptions are that the error is additive, serially, and mutually uncorrelated with the true variable and with other variables in the model (see Hausman, 2001).

precision of the prediction.

2.2 Alternative approaches to measuring monetary policy shocks

There are three main alternative approaches to measure monetary policy shocks in high-frequency and heteroskedasticity-based identification schemes.

Assuming that the data are generated by a VAR, a sufficient condition to recover the underlying shocks is that the VAR is invertible (see Stock and Watson, 2018). A necessary condition for this to hold is that the number of structural shocks in the VAR is equal to the number of variables. Otherwise, it is not possible to exactly recover the underlying structural shocks. This is a restrictive assumption, as it implies that all relevant variables are included in the VAR (see Fernández-Villaverde et al., 2007). The two-step approach does not exactly recover the true underlying shocks either, precisely because we do not assume invertibility. However, if the number of shocks equals the number of variables, the Kalman-filter prediction also exactly recovers the true shocks. Take the example in the previous subsection and assume that $R = 0$ so that the number of structural shocks is equal to the number of variables ($E = N = 1$). Then $MSE_{t|T} = \sigma_\varepsilon^2(1 - \psi^2\sigma_\varepsilon^2/(\psi^2\sigma_\varepsilon^2)) = 0$. This shows that the prediction error variance is zero, so that we can perfectly recover the unobserved shock.

Bu et al. (2021) propose using Fama-McBeth-type regressions to extract the unobserved monetary policy shocks in heteroskedasticity-based identification schemes. For every period in the sample, they regress y_t on an estimate of Ψ. However, their approach requires an orthogonality condition between the impact vector of the policy event shocks (Ψ) and other shocks (Γ), which may or may not be fulfilled in practice (see Appendix B). Meanwhile, our approach yields the best predictor in an MSE-sense without an orthogonality assumption between the impact matrices.

Finally, Bauer and Swanson (2023) propose to orthogonalize high-frequency surprises using known information at the time of the announcement. Their approach controls for the so-called 'Fed-response-to-news effect'. Our framework allows us to include relevant predictors when estimating the impact matrix Ψ using IV in the first step. The second step is based on the residual u_t, which is also orthogonal to the information set. Therefore, we readily account for predictable confounders for estimating Ψ and predicting ε_t. However, the additional advantage of the two-step approach is that we remove unpredictable confounders, such as classical measurement error, that may contaminate high-frequency surprises.

3 Simulation study

We perform a simulation study to show that the two-step predictions are more accurate than noisy proxies of event policy shocks. For this, we use estimates of Ψ, Σ, $\Gamma\Sigma_v\Gamma'$ based on data described in the next section. Specifically, we include daily changes in short-, medium-, and long-term interest rates in y_t, as well as stock prices, an exchange rate, a measure of stock market volatility, and corporate bond prices of various maturities. In addition, we control for one lag in these variables when estimating the parameters and computing the covariances. All shocks are simulated from normal distributions. We assume that there is one shock from policy events ($E = 1$) and $R = N$ other shocks. We have 8,000 observations to estimate the parameters, and a policy event occurs after 25 working days. In addition, we simulate a noisy proxy of the true underlying policy event shocks, setting the signal-to-noise ratio to 0.35. We choose this ratio as a baseline, as the correlation between the two-step prediction and the proxy amounts to 0.6, which we obtain when using actual data.

We then compute the correlation of three alternative monetary policy shock measures with the true underlying shocks: a two-step prediction based on heteroskedasticity IV, a two-step prediction based on proxy-IV, and the raw proxy. The proxy represents a noisy measure of true monetary policy shocks, similar to high-frequency surprises. Heteroskedasticity-IV estimates Ψ up to a scale using the simulated medium-term interest rate to construct a heteroskedasticity-based instrument (see Rigobon and Sack, 2004, Lewis, 2022, and Appendix A). Proxy-IV estimates Ψ up to a scale using the proxy as an instrument and the simulated medium-term interest rate as the endogenous variable (see Stock and Watson, 2018, and Appendix A).

We then vary various aspects of the simulated data to investigate under what circumstances the two-step predictions are more accurate than the raw proxy (see Figure 1). Panel a) varies the variance due to policy event shocks relative to the total variance of the medium-term interest rate. In doing so, we assume that the signal-to-noise ratio in the proxy remains constant. If policy event shocks are relatively irrelevant, the noisy proxy is a more accurate measure of the underlying shocks. However, as the variance of the policy event shocks increases, the two-step predictions become more accurate. For our baseline calibration, the two-step predictions are more accurate if more than 30% of the variance in the medium-term interest rate is driven by policy event shocks on event days.

Panel b) shows that, intuitively, the proxy performs better if it is measured with little error. However, for a signal-to-noise ratio below 0.8, the two-step predictions are more accurate in the baseline calibration. Interestingly, the signal-to-noise ratio does not affect the two-step predictions. The only exception is when the signal-to-noise-ratio is virtually zero (we set the lowest value to 0.01%). This suggests that the two-step approach based on proxy-IV works well even with very noisy instruments.

Panel c) shows results varying the number of observations to estimate the model parameters. If the model parameters are inaccurately estimated, this may increase the prediction error. However, for a reasonable number of observations, which is typically large for daily financial market data, this does not appear to be much of an issue. If the number of observations is larger around 2,000, the estimates are accurate enough to clearly outperform the noisy proxy. Panel d) shows that the number of events does not play much of a role either. Even if a policy event only occurs every quarter (roughly after 65 working days), the two-step predictions outperform the noisy proxy in the baseline calibration.

Finally, panel e) shows that more predictors tend to increase the accuracy of the predicted shocks.[8] The accuracy increases gradually with the number of variables.

The Kalman-filter prediction is only optimal under quadratic loss with normally distributed

---

[8] We reduce the number of cross-sectional responses to extract the shocks, but include all variables with one lag as control variables.

Fig. 1: Correlation of two-step predictions with true shocks

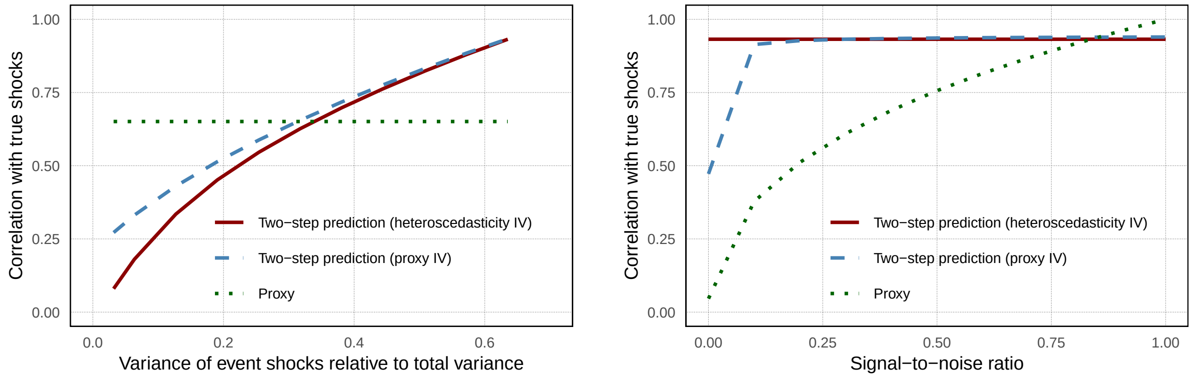

(a) Variance of event shocks

(b) Accuracy of proxy

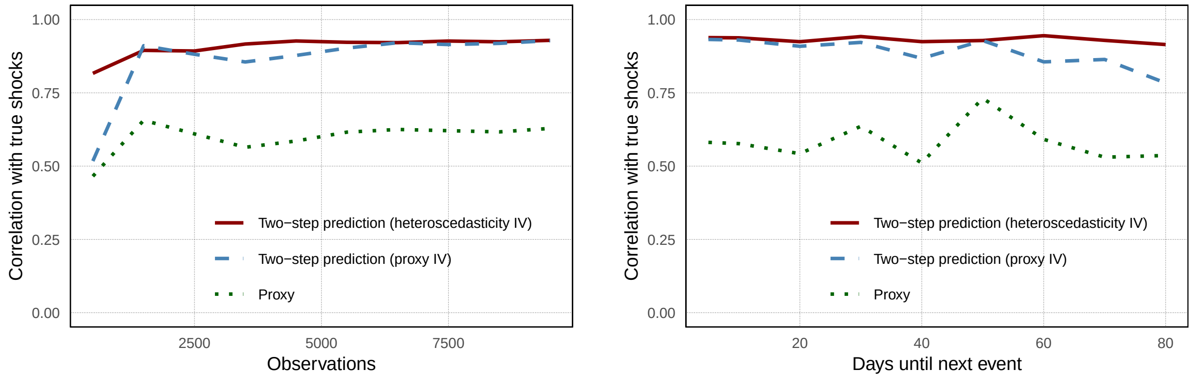

(c) Accuracy of estimates

(d) Frequency of events

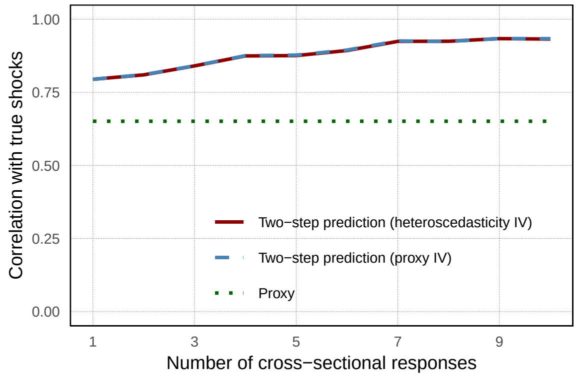

(e) Number of predictors

Notes: Correlations of monetary policy shock measures with the true underlying shocks based on simulated data. Panel a) varies the variance of event shocks relative to the total variance in the simulated medium-term interest rate. Panel b) varies the signal-to-noise ratio of the proxy of the event policy shocks. Panel c) varies the number of observations to estimate the model parameters. Panel d) varies the frequency of events. Panel e) varies the number of impact responses used to predict the unobserved policy event shocks.

shocks. Otherwise, it is only the best linear predictor. In addition, we assume in Model (1) that the variance-covariance of the underlying structural shocks is constant. Both assumptions may not be fulfilled in practice. We therefore test the robustness of the results using shock processes based on fat-tailed distributions and with conditional heteroskedasticity (see Appendix D). Qualitatively, the results for the correlation with the true underlying shocks do not change.

4 Empirical applications

We conduct various empirical applications to demonstrate the advantages of the two-step approach. First, we show that high-frequency and heteroskedasticity-based identification schemes identify similar shocks. Second, we show that some anomalies in high-frequency surprises, like the weak-instrument problem, can be alleviated by using the predicted shocks rather than the raw surprises. Third, the two-step procedure can be used to assess unexpected changes in the monetary policy stance in real-time.

4.1 Baseline specification

The baseline model we use in all applications reads:

$$\begin{aligned} \Delta y_t &= \alpha + \Psi \varepsilon_t + \Gamma v_t + \Phi(L)\Delta x_{t-1} \quad \text{for } t \in P \\ \Delta y_t &= \alpha + \Gamma v_t + \Phi(L)\Delta x_{t-1} \quad \text{for } t \in C \ , \end{aligned} \tag{2}$$

where Δ denotes the first of a variable. We set y_t to:[9]

$$
\begin{aligned}
y_t \;=\; & [i_t^{\mathrm{st}},\ i_t^{\mathrm{mt}},\ i_t^{\mathrm{lt}}, 100\times\log(\mathrm{Corp}_t),\ 100\times\log(\mathrm{Corp3Y}_t),\ 100\times\log(\mathrm{Corp5Y}_t),\\
& 100\times\log(\mathrm{Corp10Y}_t),\ 100\times\log(\mathrm{Corp15Y}_t),\\
& 100\times\log(\mathrm{NEER}_t),\ 100\times\log(\mathrm{Stocks}_t),\ \mathrm{VIX}_t,\ \mathrm{Spread}_t]'\ .
\end{aligned}
$$

Recall that the two-step approach exploits the impact matrix on financial market variables to predict the underlying shocks. Therefore, we choose variables that are likely to respond to monetary policy announcements but also display some heterogeneity in their responses. Following Bu et al. (2021), we include interest rates along the term structure, that is, a short-term (i_t^{st}), medium-term (i_t^{mt}) and a long-term interest rate (i_t^{lt}). Then we also include an overall corporate bond price index (Corp_t), subindices for various maturities, and a corporate bond spread (Spread_t). Boehm and Kroner (2024) show that stock prices and exchange rates also respond strongly to monetary policy announcements. We therefore include a nominal effective exchange rate (NEER_t), a stock price index (Stocks_t), and the VIX (VIX_t). The bond prices, nominal effective exchange rate, and stock prices are included as log multiplied by 100. Therefore, the first difference corresponds approximately to the growth rate in percent. The exchange rate is defined as one unit of foreign currency measured in terms of US Dollars (a decline corresponds to an appreciation of the US Dollar).

---

[9] A detailed description of the data sources, the exact definition and construction of the series are given in Appendix C.

As control variables, we use:

$$x_t \quad = \quad [i_t^{\text{st}},\ i_t^{\text{mt}},\ 100 \times \log(\text{NEER}_t),\ 100 \times \log(\text{Stocks}_t),\ \ 100 \times \log(\text{Spread}_t),\ \ \text{Skew}_t]' \ .$$

These control variables account for a 'Fed-response-to-news effect' (Bauer and Swanson, 2023, 2022).[10] However, compared to y_t we drop the long-term interest rate and most maturities of the corporate bond prices to avoid having to estimate too many parameters for variables that are highly correlated with each other. An additional variable we include is the implied treasury yield skewness (Skew_t) by Bauer and Chernov (2024). We do so for two reasons. First, this variable has been shown to predict high-frequency surprises (see Bauer and Swanson, 2022). Second, because it is related to upside and downside risks of treasury yields over the coming months, it may capture predictable heteroskedasticity in the residuals. We do not include this variable in y_t, because it did not respond on impact to monetary policy shocks. We set the lag length $L = 1$ in the baseline specification. This is a somewhat arbitrary choice, and we vary the lag length in a series of robustness tests.

To identify Ψ up to a scale, we resort to heteroskedasticity- and high-frequency-based instruments, assuming that there is only one monetary policy shock ($E = 1$).[11] The former approach requires policy event and control days. The monetary policy events (P), are FOMC announcement dates for 1988–2019 from Bauer and Swanson (2022). We use all

---

[10] We only include variables that are available at a daily frequency and cannot include, for example, economic news surprises used by Bauer and Swanson (2022).

[11] Therefore, we can use IV to identify the impact vector and do not have to resort to GMM as in Rigobon and Sack (2004) or use multiple high-frequency instruments as in Swanson (2021) or Altavilla et al. (2019). We leave the prediction of multi-dimensional monetary policy shocks for future research.

announcements as policy events and all other weekdays, excluding holidays, as control days. The high-frequency identification scheme is based on high-frequency surprises by Bauer and Swanson (2022). We provide results for both the orthogonalized and the raw surprises.

Tab. 1: Weak-instrument tests daily data

	(a) Heteroskedasticity		
	Short-term rate	Medium-term rate	Long-term rate
F-statistic	367.4	128.6	53.2
Effective F-statistic	144.3	40.9	11.5

	(b) High-frequency, orthogonalized		
	Short-term rate	Medium-term rate	Long-term rate
F-statistic	14.7	107.8	34.6
Effective F-statistic	5.4	62.8	22.2

	(c) High-frequency, raw		
	Short-term rate	Medium-term rate	Long-term rate
F-statistic	33.8	133.1	40.6
Effective F-statistic	9.8	68.0	26.2

Notes: The table reports the standard F-statistic, which is commonly compared to a critical value of 10 (Stock and Yogo, 2005), and the effective F-statistic, allowing for heteroskedasticity and autocorrelation, where the critical value is 23.1 (Montiel Olea and Pflueger, 2013, Lewis, 2022, Lewis and Mertens, 2022). We vary the endogenous variable and instrument along the term structure of interest rates. We show results for an average of short-term, medium-term, and long-term interest rates. The heteroskedasticity-based instruments are constructed with the corresponding endogenous variable. For the high-frequency surprises, the one-dimensional instrument remains the same, while we vary the endogenous variable. The high-frequency surprises stem from Bauer and Swanson (2022). The significance level is set to 5% and the tolerance level to 10%.

In the heteroskedasticity-based identification scheme, the question arises of which variable is used to construct the instrument (see Appendix A). Both identification schemes require specifying the endogenous variable, for which the impact effect is normalized to a specific value (unit-effect normalization).[12] If the instrument is weak for the variable on which we impose the normalization, it is well known that the impulse responses will be biased and confidence intervals unreliable.[13]

Table 1, therefore, provides weak-instrument tests for various specifications to guide this decision. We report the standard F-statistic, which is commonly compared to a critical value of 10 (Stock and Yogo, 2005), and the effective F-statistic, allowing for heteroskedasticity and autocorrelation, where the critical value is 23.1 (Montiel Olea and Pflueger, 2013, Lewis and Mertens, 2022). We prefer the latter, because it is appropriate for heteroskedasticity-based identification schemes (see Lewis, 2022), and homoskedastic errors are a strong assumption for daily financial market variables.

We vary the endogenous variable and instrument along the term structure of interest rates. We show results for an average of short-term, medium-term, and long-term interest rates. The heteroskedasticity-based instruments are constructed with the corresponding endogenous variable. For heteroskedasticity-based identification, the strongest instrument arises when the short-term interest rate is used. However, the specification based on the medium-term interest rate also passes the weak-instrument test. For the long-term interest rate, we do not reject the null hypothesis of a weak instrument based on the effective F-statistic. For high-frequency approaches, we reject weak instruments for all specifications based on the standard F-statistic.

---

[12] The first and second stages of the IV-estimation read: $\Delta i_t = \beta_0 + \beta_1 Z_t + \Theta \Delta x_{t-1} + e_t$ and $\Delta y_{it} = \alpha_i + \psi_i \Delta \hat{i}_t + \Phi_i \Delta x_{t-1} + u_{it}$, where Z_t is either the heteroskedasticity- or high-frequency-based instrument, i_t is an interest rate that we treat as the endogenous variable, e_t is a mutually and serially uncorrelated error term, and $\Delta \hat{i}_t$ is the first-stage prediction of the first difference of the interest rate.

[13] If this is the case, one could compute confidence intervals that do not depend on the instrument strength (see Montiel Olea et al., 2021).

However, using the preferred effective F-statistic, we reject weak instruments only for the medium-term interest rate and the long-term interest rate (raw surprises)

The medium-term interest rate is the only variable for which we reject weak instruments based on effective F-statistics in all specifications. In what follows, we therefore use it to construct the heteroskedasticity-based instrument and use it as the endogenous variable on which we normalize the impact effect.

4.2 High-frequency- vs. heteroskedasticity-based identification

Brennan et al. (2024) and Bu et al. (2021) find important differences between various high-frequency surprises and heteroskedasticity-based shocks. This section examines to what extent these differences stem from intraday noise in the proxies, which bias correlations towards zero.

Before turning to the predicted shocks, we investigate whether the unobserved underlying shocks have similar economic effects. We estimate daily dynamic causal effects in a long-difference local-projection framework (see Jordà, 2005).[14] For inference, we follow Montiel Olea et al. (2025) and use heteroskedasticity-consistent (HC) standard errors.

Figure 2 shows local-projection estimates based on heteroskedasticity (solid blue line) and high-frequency surprises (dotted red line). Panels (a) and (b) vary between orthogonalized and raw surprises, respectively. The underlying monetary policy shocks have similar causal effects

---

[14] Compared to a level specification, this framework is more robust when estimating persistent responses (see Piger and Stockwell, 2025). Compared to a SVAR, it is more robust with respect to model-misspecification (see Montiel Olea et al., 2024). See Montiel Olea et al. (2025) for an overview of the advantages and disadvantages compared to VARs.

on key financial variables. The point estimates from the high-frequency schemes are mostly statistically indistinguishable from the heteroskedasticity-based estimates. The only exception are stock prices, which show a somewhat stronger response after 15 working days using the high-frequency schemes.

Fig. 2: Daily dynamic causal effects

(a) Heteroskedasticity vs. orthogonalized high-frequency surprises

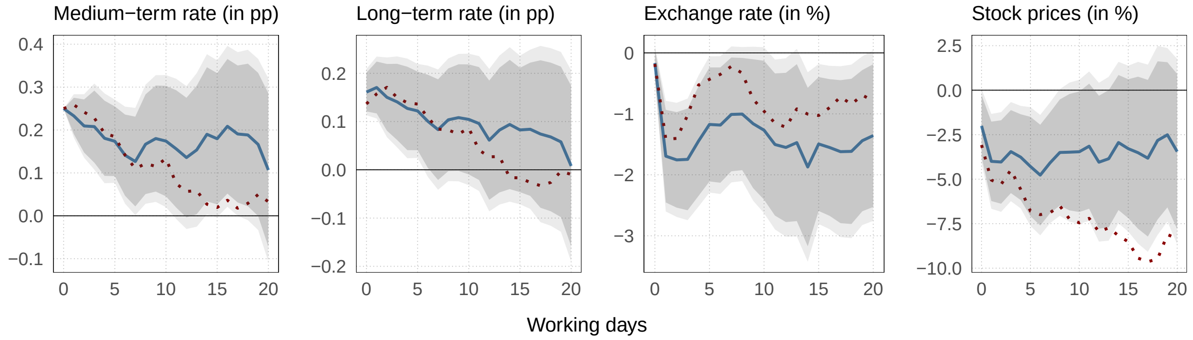

(b) Heteroskedasticity vs. raw high-frequency surprises

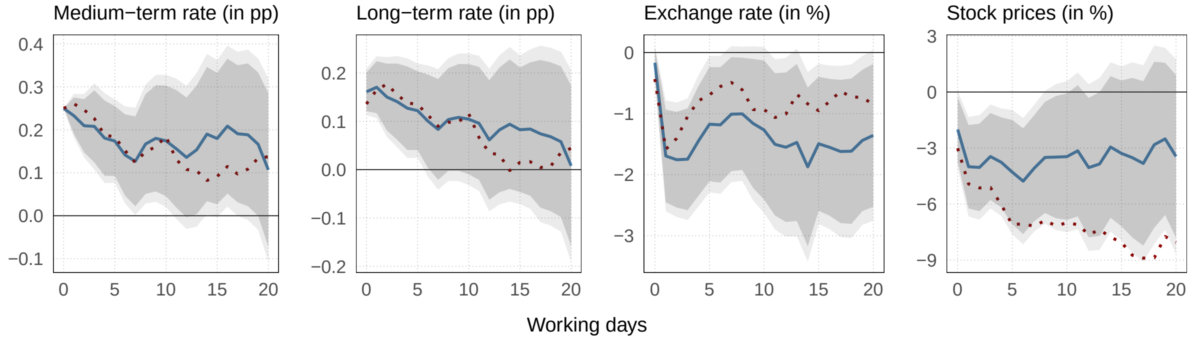

Notes: Impulse responses to monetary policy shocks. The responses are normalized to a 25 bp increase in the medium-term interest rate. The solid blue lines show the heteroskedasticity-based responses estimated using IV. The red dotted lines are the responses to the high-frequency surprises by Bauer and Swanson (2022) estimated using IV. The horizontal axis measures working days (excluding weekends and holidays). The models are estimated in first (log-)differences, but the impulse responses are cumulated. Therefore, all interest rate responses are measured in percentage points and the exchange rate responses are measured in percent. The gray areas show 90% and 95% confidence intervals, which are based on HC standard errors.

In addition to examining the dynamic responses, we can also investigate the impact vector, that is, the cross-sectional variation of the responses of the financial variables (see Table 2). The two identification schemes yield qualitatively similar results. A 0.25 pp increase in the medium-term interest rate, caused by a monetary policy shock, leads to a somewhat smaller increase in the short-term interest rate and an even smaller increase in the long-term interest rate. Meanwhile, the Dollar appreciates, stock prices fall, and stock price volatility increases. The shock is also transmitted to corporate bond prices that fall across the maturity spectrum. Meanwhile, the corporate bond spread declines. Most of the impact effects are statistically significant, at least at the 10% level. The only exception is the impact effect on the exchange rate, which is not significant for the heteroskedasticity-based scheme and when using the orthogonalized surprises.[15]

Tab. 2: Impact vector estimates

	Heteroskedasticity	High-frequency, orthogonalized	High-frequency, raw
Medium-term rate (in pp)	0.25***	0.25***	0.25***
Short-term rate (in pp)	0.20***	0.15***	0.21***
Long-term rate (in pp)	0.16***	0.14***	0.14***
Exchange rate (in %)	-0.17	-0.18	-0.43***
Stock prices (in %)	-2.01*	-3.08***	-3.01***
VIX (in pp)	2.15*	2.33***	2.11***
Corporate bond spread (in pp)	-0.13***	-0.08***	-0.08***
Corp. bond index (in %)	-0.75***	-0.86***	-0.79***
Corp. bond index 3Y (in %)	-0.42***	-0.49***	-0.46***
Corp. bond index 5Y (in %)	-0.74***	-0.82***	-0.75***
Corp. bond index 10Y (in %)	-0.87***	-1.04***	-0.95***
Corp. bond index 15Y (in %)	-0.60**	-0.80***	-0.79***

Notes: IV estimates of the impact vectors using heteroskedasticity- and high-frequency-based instruments. The orthogonalized and raw high-frequency surprises stem from Bauer and Swanson (2022). Inference is based on HC standard errors. ***/**/* denotes statistical significance at the 1%/5%/10% level.

The dynamic and cross-sectional effects suggest that we identify similar underlying monetary

---

[15] This may be due to the fact that the exchange rate data is recorded at noon EST, whereas most other series are recorded at market close.

policy shocks. According to Proposition (1), we can use the impact matrix, along with the residuals and the variance-covariance matrix of the model, to predict these shocks up to a scale and compare them across identification strategies. Table 3 suggests that the low correlation between various shock measures reported in the literature is related to the fact that they are measured with error. The upper panel shows the correlations for the two-step predictions based on various identification schemes. The correlations are typically 0.82 and higher. The two-step predictions based on the raw surprises and the orthogononalized surprises even show a correlation of 0.97; if we compare the two-step predictions to the high-frequency surprises themselves, they are clearly lower (between 0.53 and 0.70).

Tab. 3: Correlations between monetary policy shocks and high-frequency surprises

	Two-step predictions		
	Heteroskedasticity	High-frequency, orthogonalized	High-frequency, raw
Two-step predictions			
Heteroskedasticity	1.00		
High-frequency, orth.	0.82	1.00	
High-frequency, raw	0.83	0.97	1.00
Surprises			
Orthogonalized	0.53	0.66	0.64
Raw	0.57	0.67	0.70

Notes: Correlations between two-step predictions of monetary policy shocks based on heteroskedasticity and high-frequency identification schemes, as well as the monetary policy surprises from Bauer and Swanson (2022).

These findings are in line with Brennan et al. (2024) and Bu et al. (2021), who find correlations as low as 0.3 between various monetary policy shock measures based on various identification and measurement schemes, but qualitatively similar causal effects. We complement their findings by proposing that the two-step prediction yields a more accurate measure of the true underlying monetary policy shocks. However, predictions are also subject to errors. We therefore turn next to an evaluation of the two-step predictions in terms of various anomalies typically encountered when using existing monetary policy shock measures.

4.3 Anomalies in monetary policy shock measures

Weak-instrument problem The literature has documented that high-frequency surprises produce weak instruments when used to identify the causal effect of monetary policy in monthly or quarterly macroeconomic models (Ramey, 2016, Bauer and Swanson, 2022). Instruments can be weak because monetary policy has become more predictable over time, so the variance of monetary policy surprises has declined. However, if monetary policy announcements affect financial market variables with a delay, a narrow event window may not capture the full effect of the announcement. Finally, unpredictable intraday noise in the surprises will lower the strength of the instrument. To investigate this question, we estimate monthly local projections using data on US medium-term interest rates, industrial production, consumer prices, and stock prices.[16] We estimate the model in first (log-)differences and include 12 lags of all variables as controls. To identify dynamic causal effects, we examine various interest rates as endogenous variables. As instruments we use the sum of the daily monetary policy shock measures over each month as instruments.[17]

---

[16] A detailed description of the data is given in Appendix C.

[17] The monthly model largely follows Gertler and Karadi (2015) and Bauer and Swanson (2022), with some differences. We estimate local projections with differenced data. Also, we estimate the coefficients on a smaller sample, where the shocks are available starting in 1988. We also do not include the excess bond premium, to keep the model parsimonius.

Table 1 provides weak-instrument tests. Again, we report the standard and effective F-statistics, where the critical value is 10 for the former and 23.1 for the latter. For all specifications, the two-step predictions yield stronger monthly instruments than the corresponding surprises. This suggests that the two-step approach filters out unpredictable daily noise from the data. In line with Bauer and Swanson (2022), orthogonalized surprises yield weaker instruments than raw surprises. However, in both cases, the two-step predictions yield somewhat stronger instruments. The two-step predictions based on heteroskedasticity yield the strongest instruments and, for the medium-term interest rate, also pass the weak-instrument test according to the effective F-statistic.

The results show that the two-step prediction alleviates the weak-instrument problem. One explanation for this result is that high-frequency surprises include intraday noise that cannot be filtered out based on the information known before the announcement. However, the cross-sectional response of the financial variables, exploited in the two-step prediction, provides useful information on the true underlying monetary policy shock. In addition, the heteroskedasticity-based approach, which captures the entire daily variation and does not need a narrow event time window, yields the strongest instruments.

For comparison, Bauer and Swanson (2022) obtain standard F-statistics of 7.7 and 2.4 for their raw and orthogonalized surprises, respectively. When they add relevant speeches to the event data set, their F-statistics increase to 15.8 and 10.2, respectively.[18] Our approach yields stronger

---

[18] It would be interesting in future research to apply the two-step approach to the relevant speeches collected by Swanson and Jayawickrema (2023). The heteroskedasticity-based two-step approach may be particularly suited because speeches occur over longer time periods and therefore the exact time stamp of an announcement is more difficult to determine. However, at the time of writing those dates were not publicly available.

Tab. 4: Weak-instrument tests monthly data

	(a) Two-step prediction, heteroskedasticity		
	Short-term rate	Medium-term rate	Long-term rate
F-statistic	15.4	11.4	5.9
Effective F-statistic	9.0	24.3	7.9

	(b) Two-step prediction, high-frequency, orthogonalized		
	Short-term rate	Medium-term rate	Long-term rate
F-statistic	10.1	4.3	1.7
Effective F-statistic	6.9	4.1	1.6

	(c) Two-step prediction, high-frequency, raw		
	Short-term rate	Medium-term rate	Long-term rate
F-statistic	14.8	6.1	1.9
Effective F-statistic	11.2	6.6	1.7

	(d) Surprises, high-frequency, orthogonalized		
	Short-term rate	Medium-term rate	Long-term rate
F-statistic	0.2	0.3	1.3
Effective F-statistic	0.4	0.2	1.9

	(e) Surprises, high-frequency, raw		
	Short-term rate	Medium-term rate	Long-term rate
F-statistic	9.5	2.8	0.2
Effective F-statistic	7.5	1.5	0.1

Notes: The table reports the standard F-statistic, which is commonly compared to a critical value of 10 (Stock and Yogo, 2005), and the effective F-statistic, allowing for heteroskedasticity and autocorrelation, where the critical value is 23.1 (Montiel Olea and Pflueger, 2013, Lewis, 2022, Lewis and Mertens, 2022). We vary the endogenous variable and instrument along the term structure of interest rates. We show results for an average of short-term, medium-term and long-term interest rates. As instruments, we use the monthly sum of the two-step predictions and the monetary policy surprises by Bauer and Swanson (2022). The significance level is set to 5% and the tolerance level to 10%.

instruments, in a similar order of magnitude, without adding additional events.

The estimates of the monthly dynamic causal effects are also in line with economic intuition (see Figure 3). we find qualitatively similar results when using the two-step predictions based on the heteroskedasticity- and high-frequency-based identification schemes. A 0.25 pp increase in the medium-term interest rate, caused by a monetary policy shock, leads to a decline in industrial production, a decline of the CPI and a drop in stock prices.

Fig. 3: Impulse responses monthly macroeconomic variables

(a) Two-step, heteroskedasticity

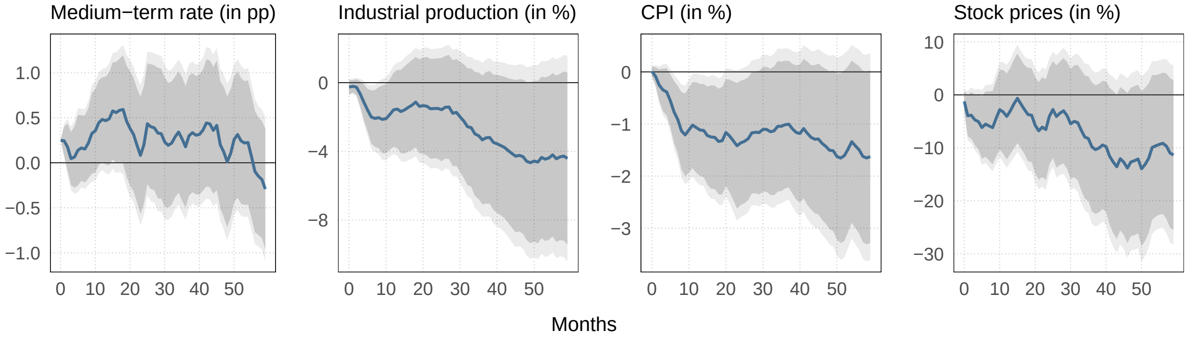

(b) Two-step, high-frequency, raw

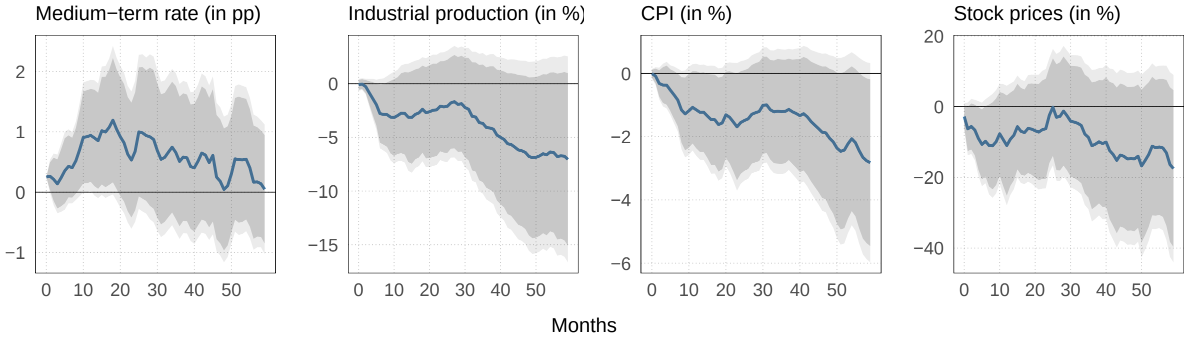

Notes: Impulse responses to monetary policy shocks. The responses are normalized to a 25 bp increase in the medium-term interest rate caused by a monetary policy shock. The horizontal axis measures months. The models are estimated in first (log-)differences, but the impulse responses are cumulated. Therefore, all interest rate responses are measured in percentage points and all other responses are measured in percent. The gray areas show 90% and 95% confidence intervals, which are based on HC standard errors.

'Fed-response-to news effect' We next turn to the observation that high-frequency surprises are predictable based on information available before the announcement (see Bauer and Swanson, 2023, 2022). The two-step prediction also removes predictable movements in monetary policy shocks because it accounts for information available at the time of the announcement.

Tab. 5: Predictability of monetary policy shocks and surprises

	(1)	(2)	(3)	(4)
Nonfarm payrolls	0.17**	0.13**	0.11	
	(0.07)	(0.06)	(0.06)	
Empl. growth (12m)	0.13**	0.14***	0.11*	0.14**
	(0.06)	(0.05)	(0.06)	(0.06)
Log-diff. S&P 500 (3m)	0.10	−0.02	0.04	
	(0.07)	(0.07)	(0.08)	
Diff. slope (3m)	−0.08	−0.01	−0.02	
	(0.05)	(0.06)	(0.07)	
Log-diff. comm. price (3m)	0.16**	0.13	0.03	
	(0.07)	(0.08)	(0.08)	
Treasury skewness	0.15***	0.12**	0.05	
	(0.05)	(0.06)	(0.06)	
Measurement	Surprises	Two-step	Two-step	Two-step
Identification	High-freq., raw	High-freq., orth.	Heterosc.	Heterosc.
Observations	315	315	315	315
R^2	0.17	0.08	0.04	0.02

Notes: Regressions of high-frequency surprises (1) and two-step predictions of monetary policy shocks (2-4) on information known before the announcement. The high-frequency surprises and predictor variables stem from Bauer and Swanson (2022). Inference is based on HC standard errors. ***/**/* denotes statistical significance at the 1%/5%/10% level.

To show that the two-step predictions are less predictable than the raw high-frequency surprises, we replicate the regression used by Bauer and Swanson (2022) to orthogonalize the shocks in Table 5. The predictor variables available before the announcement are the surprise component of the last nonfarm payroll release, the log-change in nonfarm payrolls over the past 12 months, the log change in the S&P 500 over the last three months, the change in the yield curve over the last three months, the log-change in a commodity price index over the last 3 months, and the implied skewness of the treasury yield on the day before the announcement.

Column (1) regresses the raw high-frequency surprises on the predictors. Surprises in nonfarm payrolls, employment growth, commodity price growth, and treasury yield skewness exhibit a statistically significant coefficient. The predictors jointly explain 17% of the variation in high-frequency surprises. Column (2) shows the same regression using the two-step prediction based on orthogonalized surprises. Recall that those account for one lag of the control variables included in our baseline specification. Three predictors are still statistically significant, but the absolute size of most coefficients is smaller. As a result, the regressors jointly explain only 8% of the variation of the two-step prediction. The explanatory power drops even further to 4% when using the heteroskedasticity-based two-step prediction. Only one predictor, employment growth over the past 12 monthy, shows a significant coefficient at the 10% level. Finally, removing predictors that are not statistically significant in column (4), the explanatory power falls to to 2%. This implies that, although the monetary policy shock measure is still predictable, the predictable part makes up only a minor share of the total variance. These results suggest that the two-step prediction framework not only accounts for unpredictable noise in the data, but also alleviates the 'Fed-response-to-news effect'.

Information effect The final anomaly that we investigate is the information effect. Researchers have found that stock prices move in the same direction as monetary policy surprises for some announcements (see Nakamura and Steinsson, 2018, Jarociński and Karadi, 2020). An explanation for this pattern is an information effect. That is, by increasing the short-term interest rate, the central bank may not only surprise market participants because they implement a more contractionary monetary policy for a given economic situation (a monetary policy shock implying negative co-movement) but also convey to markets that the economic situation is better than they anticipated, which warrants an endogenous response by the central bank (an information shock implying a positive co-movement).

Jarociński and Karadi (2020) use sign restrictions in a proxy-VAR to disentangle the two shocks. But also, they use 'poor-man's sign restrictions' in a robustness test. That is, to identify a monetary policy (information) shock, they only use high-frequency surprises for which the sign is the opposite (the same) as stock price movements. However, Bauer and Swanson (2023) have questioned the relevance of the information effect, suggesting that it may be an artifact of not controlling the information that the market participants had before the actual announcement. In addition, if high-frequency surprises are subject to measurement error, there will be additional anomalous co-movements, which are completely spurious.

Tab. 6: Share of periods with positive co-movement with stock prices

Two-step predictions			Surprises	
Heteroskedasticity	High-frequency, orthogonalized	High-frequency, raw	High-frequency, orthogonalized	High-frequency, raw
0.33	0.29	0.31	0.42	0.44

Notes: Share of monetary policy shocks and surprises having the same sign as stock price returns on the same day.

To investigate the relative importance of predictable information and unpredictable noise to explain the anomalous positive co-movement between monetary policy shock measures and stock prices, we compute the share of announcements with positive daily co-movement (see Table 6).[19] A lower share implies that more announcements display the expected negative co-movement. Unpredictable noise is more important to explain anomalous co-movements between monetary policy shock measures and stock prices than predictable information. The raw high-frequency surprises move in the same direction in 44% of the announcements. When using orthogonalized surprises, this share falls to 42%. When using the two-step predictions, the share declines to between 29% and 33%, depending on the specification. This suggests that the share of positive co-movement between monetary policy shocks and stock prices is up to 13 pp lower when we account for unpredictable intraday noise using the two-step approach.

4.4 Real-time assessment of monetary policy announcements

A key question for policymakers is whether an announcement had a surprising contractionary or expansionary effect. We therefore investigate whether the two-step approach can predict monetary policy shocks in real time. The predicted shocks may suffer from revisions if the model parameters change as additional data becomes available. In what follows, we examine the accuracy of the real-time shock predictions based on a pseudo-out-of-sample forecasting exercise.

The real-time evaluation begins with the first policy announcement in 2000 to ensure a sufficient number of initial observations. From this point on, we predict vintages of the shocks

---

[19] We experimented with various stock price indices. The results are qualitatively similar.

based on models estimated using data available one day after each policy decision. The initial real-time estimate of the shocks then corresponds to the last value of each vintage. We compare this initial estimate to the final estimate, that is, the last vintage estimated using the entire sample. Then, we also examine whether the initial estimate is revised at the next two meetings, by re-estimating the model based on the additional information and predict the shocks based on the updated parameters.

Tab. 7: Accuracy of real-time predictions

	Correlation		
	First prediction	Second prediction	Third prediction
Heteroskedasticity	0.79	0.80	0.80
High-frequency, raw	0.93	0.96	0.96
High-frequency, orth.	0.92	0.96	0.96
	Share correct sign		
	First prediction	Second prediction	Third prediction
Heteroskedasticity	0.85	0.85	0.85
High-frequency, raw	0.95	0.94	0.93
High-frequency, orth.	0.95	0.95	0.95

Notes: The table shows correlations and the share of correct signs of the two-step predictions of monetary policy shocks. Each row uses a different instrument to estimate the impact matrix. Each column shows the accuracy of the first, second prediction, and third predictions, compared to the prediction using the entire sample. The state of information of the first prediction is data available one day after the announcement. For the second (third) prediction the state of information is what has been available one day after the next announcement (the announcement after the next).

The results in Table 7 show that the initial estimates are good predictors of the final estimate. The correlation between the initial and final estimates is greater than 0.9 for the two-step predictions based on high-frequency identification. For the heteroskedasticity-based two-step predictions, the correlation is lower but still around 0.8.

A central bank may also be interested in whether the announcement has been surprisingly expansionary or contractionary. We therefore report the share of announcements where the real-time predictions have the same sign as the final estimate in the second panel. In the vast majority of cases, the first prediction correctly determines the sign of the underlying shock. The share is 95% for high-frequency-based identification and 0.85% for heteroskedasticity-based identification.

4.5 Robustness tests

The predictions of daily monetary policy shocks are relatively robust with respect to changing the model specification (see Appendix D).

Varying the lag length of the control variables leads to virtually identical results.

We then estimate the model on a longer sample from 1986 – 2022. To do so, we collect our own events before and after the sample of Bauer and Swanson (2022). Prior to 1988, we collected dates of discount rate changes and Federal Funds target changes. After 2019, we added FOMC announcements. The data sources are given in Appendix C. In total, we added 40 additional events to the data set. The correlation is somewhat lower, but still high (0.94). Note that this may stem primarily from the fact that we have to exclude variables with missing data prior to 1988 for shock prediction (five bond price indices) and from the control data set (treasury yield skewness).

We then exclude events that may change the variance of financial market variables on control days but not on policy event days. Specifically, we collect dates of FOMC minutes releases, speeches by the Chair and Vice Chair of the Board of Governors, decisions by other central banks, and important data releases. Note that some of those events start after 1988, for example, some data release dates and central bank decisions (see Appendix C).[20] But the model is still estimated on the baseline sample. The correlation with the baseline remains very high (0.99).

5 Concluding remarks

We propose a two-step approach to measure monetary policy shocks based on daily financial market data. The first step uses standard IV approaches to identify the causal effect of monetary policy shocks on various daily financial market variables. The second step predicts the underlying unobserved shocks from the cross-sectional variation in the responses of financial variables.

The two-step approach has various desirable properties. It delivers the optimal (linear) prediction under quadratic loss. It uses readily available data and identification schemes. Under realistic assumptions on the data generating process and sample size, it delivers more accurate measurements of the underlying monetary policy shocks compared to noisy proxies of the shocks. This implies that various anomalies observed in the literature on high-frequency surprises are alleviated or resolved. Finally, for policy makers, it is important to measure monetary policy shocks in real-time to understand unexpected changes in the policy stance. We show that the two-step approach delivers accurate predictions in real time.

---

[20] In the working paper version, we also identified relevant speeches using a topic-modeling approach (see Burri and Kaufmann, 2024). Excluding relevant speeches did not change the results either.

Our approach can also be applied in other situations. It would be interesting to predict multi-dimension shocks, distinguishing between target, forward guidance, and large-scale asset purchase shocks.[21] Additionally, it would be interesting to apply the approach if monetary policy events are known, but high-frequency instruments are missing due to a lack of intraday tick data or lack of knowledge of the exact intraday timing of announcements. This could be especially relevant for examining the role of speeches or monetary policy shocks on longer sample periods. Finally, researchers also use daily financial market data in high-frequency identification schemes (see Bolhuis et al., 2024, An et al., 2025). However, these daily changes may also contain more intraday noise such that the two-step prediction would be particularly useful for refining these measurements. We leave these extensions for future research.

---

[21] For example, there exist multi-dimensional monetary policy surprises based on high-frequency data (see Altavilla et al., 2019, Swanson, 2021).

References

Altavilla, C., Brugnolini, L., Gürkaynak, R. S., Motto, R., and Ragusa, G. (2019). Measuring euro area monetary policy. *Journal of Monetary Economics*, 108:162–179, DOI: 10.1016/j.jmoneco.2019.08.016.

An, P., Stedman, K. D., and Lusompa, A. (2025). How high does high frequency need to be? a comparison of daily and intradaily monetary policy surprises. Working Paper 25-03, Federal Reserve Bank of Kansas City.

Bauer, M. and Chernov, M. (2024). Interest rate skewness and biased beliefs. *The Journal of Finance*, 79(1):173–217, DOI: 10.1111/jofi.13276.

Bauer, M. D. and Swanson, E. T. (2022). A reassessment of monetary policy surprises and high-frequency identification. In Eichenbaum, M., Hurst, E., and Ramey, V. A., editors, *NBER Macroeconomics Annual 2022, volume 37*. University of Chicago Press.

Bauer, M. D. and Swanson, E. T. (2023). An alternative explanation for the “Fed information effect”. *American Economic Review*, 113(3):664–700, DOI: 10.1257/aer.20201220.

Boehm, C. and Kroner, T. N. (2024). Monetary policy without moving interest rates: The Fed non-yield shock. Working Paper 32636, National Bureau of Economic Research.

Bolhuis, M. A., Das, S., and Yao, B. (2024). A new dataset of high-frequency monetary policy shocks. Working Paper 224, International Monetary Fund, DOI: 10.5089/9798400290930.001.

Braun, R., Miranda-Agrippino, S., and Saha, T. (2024). Measuring monetary policy in the UK: The UK monetary policy event-study database. *Journal of Monetary Economics*, DOI: 10.1016/j.jmoneco.2024.103645.

Brennan, C. M., Jacobson, M. M., Matthes, C., and Walker, T. B. (2024). Monetary policy shocks: Data or methods? Finance and Economics Discussion Series 2024-001, Federal Reserve Board, DOI: 10.17016/FEDS.2024.011.

Bu, C., Rogers, J., and Wu, W. (2021). A unified measure of Fed monetary policy shocks. *Journal of Monetary Economics*, 118:331–349, DOI: 10.1016/j.jmoneco.2020.11.002.

Burri, M. and Kaufmann, D. (2024). Measuring multiple monetary policy shocks based on heteroscedasticity. Working Paper 24-03, IRENE Institute of Economic Research, University of Neuchâtel.

Burri, M. and Kaufmann, D. (2025). Causal effects of multiple monetary policy dimensions: A heteroskedasticity IV approach. Mimeo, IRENE Institute of Economic Research, University of Neuchâtel.

Canetg, F. and Kaufmann, D. (2022). Overnight rate and signalling effects of central bank bills. *European Economic Review*, 143:104060, DOI: 10.1016/j.euroecorev.2022.104060.

Casini, A. and McCloskey, A. (2025). Identification and estimation of causal effects in high-frequency event studies. Working Paper 2406.15667, arXiv.org, DOI: 10.48550/arXiv.2406.15667.

Fama, E. F. and MacBeth, J. D. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, 81(3):607–636, DOI: 10.1086/260061.

Fernández-Villaverde, J., Rubio-Ramírez, J. F., Sargent, T. J., and Watson, M. W. (2007). ABCs (and Ds) of understanding VARs. *American Economic Review*, 97(3):1021–1026, DOI: 10.1257/aer.97.3.1021.

Gertler, M. and Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, 7(1):44–76, DOI: 10.1257/mac.20130329.

Gürkaynak, R. S., Sack, B., and Swanson, E. (2005). Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. *International Journal of Central Banking*, 1(1).

Gürkaynak, R. S., Sack, B., and Wright, J. H. (2007). The U.S. Treasury yield curve: 1961 to the present. *Journal of Monetary Economics*, 54(8):2291–2304, DOI: 10.1016/j.jmoneco.2007.06.029.

Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press.

Hausman, J. (2001). Mismeasured variables in econometric analysis: Problems from the right and problems from the left. *Journal of Economic Perspectives*, 15(4):57–67, DOI: 10.1257/jep.15.4.57.

Jarociński, M. and Karadi, P. (2020). Deconstructing monetary policy surprises – The role of information shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43, DOI: 10.1257/mac.20180090.

Jordà, O. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182, DOI: 10.1257/0002828053828518.

Kolesár, M. and Plagborg-Møller, M. (2025). Dynamic causal effects in a nonlinear world: the good, the bad, and the ugly. *Journal of Business & Economic Statistics*, 43(4):737–754, DOI: 10.1080/07350015.2025.2539478.

Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics*, 47(3):523–544, DOI: 10.1016/S0304-3932(01)00055-1.

Lewis, D. J. (2022). Robust inference in models identified via heteroskedasticity. *The Review of Economics and Statistics*, 104(3):510–524, DOI: 10.1162/rest_a_00963.

Lewis, D. J. and Mertens, K. (2022). A robust test for weak instruments with multiple endogenous regressors. Staff Reports 1020, Federal Reserve Bank of New York.

Miranda-Agrippino, S. and Ricco, G. (2021). The transmission of monetary policy shocks. *American Economic Journal: Macroeconomics*, 13(3):74–107, DOI: 10.1257/mac.20180124.

Montiel Olea, J. L. and Pflueger, C. (2013). A robust test for weak instruments. *Journal of Business & Economic Statistics*, 31(3):358–369, DOI: 10.1080/00401706.2013.806694.

Montiel Olea, J. L., Plagborg-Møller, M., Qian, E., and Wolf, C. K. (2024). Double robustness of local projections and some unpleasant varithmetic. Working Paper 32495, National Bureau of Economic Research, DOI: 10.3386/w32495.

Montiel Olea, J. L., Plagborg-Møller, M., Qian, E., and Wolf, C. K. (2025). Local projections or VARs? A primer for macroeconomists. In *NBER Macroeconomics Annual 2025*, volume 40. University of Chicago Press.

Montiel Olea, J. L., Stock, J. H., and Watson, M. W. (2021). Inference in structural vector autoregressions identified with an external instrument. *Journal of Econometrics*, 225(1):74–87, DOI: 10.1016/j.jeconom.2020.05.014.

Nakamura, E. and Steinsson, J. (2018). High-frequency identification of monetary non-neutrality: The information effect. *The Quarterly Journal of Economics*, 133(3):1283–1330, DOI: 10.1093/qje/qjy004.

Piger, J. and Stockwell, T. (2025). Differences from differencing: Should local projections with observed shocks be estimated in levels or differences? *Journal of Applied Econometrics*, DOI: 10.1002/jae.70003.

Ramey, V. (2016). Macroeconomic shocks and their propagation. volume 2 of *Handbook of Macroeconomics*, pages 71–162. Elsevier, DOI: 10.1016/bs.hesmac.2016.03.003.

Rigobon, R. (2003). Identification through heteroskedasticity. *The Review of Economics and Statistics*, 85(4):777–792.

Rigobon, R. and Sack, B. (2004). The impact of monetary policy on asset prices. *Journal of Monetary Economics*, 51(8):1553–1575, DOI: 10.1016/j.jmoneco.2004.02.004.

Stock, J. H. and Watson, M. W. (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *The Economic Journal*, 128(610):917–948, DOI: 10.1111/ecoj.12593.

Stock, J. H. and Yogo, M. (2005). Testing for weak instruments in linear IV regression. In Andrews, D. W. K. and Stock, J. H., editors, *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, pages 80–108. Cambridge University Press, Cambridge.

Swanson, E. T. (2021). Measuring the effects of Federal Reserve forward guidance and asset purchases on financial markets. *Journal of Monetary Economics*, 118:32–53, DOI: 10.1016/j.jmoneco.2020.09.003.

Swanson, E. T. and Jayawickrema, V. (2023). Speeches by the Fed Chair are more important than FOMC announcements: An improved high-frequency measure of U.S. monetary policy shocks. *mimeo*.

Zhu, L. (2023). Let the market speak: Using interest rates to identify the Fed information effect. Working paper, SSRN, DOI: 10.2139/ssrn.4035869.

Appendix

A First step: Identification and estimation of Ψ

In what follows, we discuss the identification and estimation of the impact matrix Ψ in Model (1) using heteroskedasticity and high-frequency-based schemes. It follows existing work by Rigobon (2003), Rigobon and Sack (2004) and Stock and Watson (2018), for example, and serves mainly to keep the paper self-contained. We use Σ_ε and Σ_v to denote the variance-covariance matrices of the structural shocks and Ψ_i and Γ_i the corresponding impact vectors for variable i. Furthermore, u_i denotes a column vector of residuals $t = 1, \ldots, T$ of a regression of the variable y_i on the control variables. Finally, we denote $x_{it \in P}$ a column vector that contains only observations on policy event days, and $\mathbb{COV}[x_i, x_j]_{t \in P}$ a covariance on policy event days.

A.1 Heteroskedasticity-based IV estimator $E = 1$

Assume that $E = 1$. We can use a heteroskedasticity-based instrument Z_{1t} using the residual of the variable 1 (see Rigobon and Sack, 2004, Lewis, 2022):

$$Z_{1t} \quad = \quad \left[\mathbf{1}(t \in P)\frac{T}{T_P} - \mathbf{1}(t \in C)\frac{T}{T_C}\right] u_{1t}$$

Note that:

$$\begin{aligned}
\frac{1}{T} Z_1' u_i &= \frac{1}{T_P} u'_{1t\in P} u_{it\in P} - \frac{1}{T_C} u'_{1t\in C} u_{it\in C} \\
&\xrightarrow{p} \mathbb{COV}[u_1, u_i]_{t\in P} - \mathbb{COV}[u_1, u_i]_{t\in C} \\
&= \tilde{\mathbb{COV}}[u_1, u_i] \ ,
\end{aligned}$$

where we use the notation $\tilde{\mathbb{COV}}[u_i, u_j] = \mathbb{COV}[u_i, u_j]_{t\in P} - \mathbb{COV}[u_i, u_j]_{t\in C}$ to denote the difference in covariance on policy event and other days.

Therefore, when using Z_1 as an instrument for y_1, the IV estimator of the impact effect on the variable i reads:

$$\begin{aligned}
\Psi_{i1}^{IV} &= \frac{\frac{1}{T} Z_1' u_i}{\frac{1}{T} Z_1' u_1} \\
&\xrightarrow{p} \frac{\tilde{\mathbb{COV}}[u_1, u_i]}{\tilde{\mathbb{V}}[u_1]}
\end{aligned}$$

Using Model (1) we obtain:

$$\begin{aligned}
\tilde{\mathbb{COV}}[u_1, u_i] &= \Psi_{11}\Psi_{i1}\sigma_{1\varepsilon}^2 + \Gamma_1\Sigma_v\Gamma_i' - \Gamma_1\Sigma_v\Gamma_i' = \Psi_{11}\Psi_{i1}\sigma_{1\varepsilon}^2 \\
\tilde{\mathbb{V}}[u_1] &= \Psi_{11}^2\sigma_{1\varepsilon}^2 + \Gamma_1\Sigma_v\Gamma_1' - \Gamma_1\Sigma_v\Gamma_1' = \Psi_{11}^2\sigma_{1\varepsilon}^2
\end{aligned}$$

It follows that:

$$\Psi_{i1}^{IV} \xrightarrow{p} \frac{\Psi_{11}\Psi_{i1}\sigma_{1\varepsilon}^2}{\Psi_{11}^2\sigma_{1\varepsilon}^2} = \frac{\Psi_{i1}}{\Psi_{11}} \propto \Psi_{i1}$$

Therefore, IV consistently estimates Ψ up to a scale, assuming that $\Psi_{11} \neq 0$. The latter corresponds to the assumption that the variance of u_1 indeed increases on days with policy events compared to other days.

A.2 High-frequency-based IV estimator $E = 1$

Assume $E = 1$. Further, assume that for policy event days, we obtain a noisy measure (proxy) of shock 1:

$$Z_{1t} = \varepsilon_{1t} + \eta_{1t} \text{ for } t \in P ,$$

where η_{1t} is a serially and mutually uncorrelated random variable with zero mean and variance $\sigma_{1\eta}^2$. Because η_{1t} and v_t are serially and mutually uncorrelated shocks we have:

$$\begin{aligned} \frac{1}{T_P} Z'_{1t\in P} u_{it\in P} &= \frac{1}{T_P}\varepsilon'_{1t\in P}[\Psi_{i1}\varepsilon_{1t\in P} + \Gamma_i v_{1t\in P}] + \frac{1}{T_P}\eta'_{1t\in P}[\Psi_{i1}\varepsilon_{1t\in P} + \Gamma_i v_{t\in P}] \\ &\xrightarrow{p} \mathbb{COV}[\varepsilon_1, \Psi_{i1}\varepsilon_1]_{t\in P} \\ &= \Psi_{i1}\sigma_{1\varepsilon}^2 \end{aligned}$$

When using the instrument Z_1 for the endogenous variable y_1, the IV estimator of the impact effect on variable *i* reads:

$$\begin{aligned}
\Psi_{i1}^{IV} &= \frac{\frac{1}{T_P} Z'_{1t\in P} u_{it\in P}}{\frac{1}{T_P} Z'_{1t\in P} u_{1t\in P}} \\
&\xrightarrow{p} \frac{\mathbb{COV}[\varepsilon_1, \Psi_{i1}\varepsilon_1]_{t\in P}}{\mathbb{COV}[\varepsilon_1, \Psi_{11}\varepsilon_1]_{t\in P}} \\
&= \frac{\Psi_{i1}\sigma_{1\varepsilon}^2}{\Psi_{11}\sigma_{1\varepsilon}^2} = \frac{\Psi_{i1}}{\Psi_{11}} \propto \Psi_{i1}
\end{aligned}$$

Therefore, IV consistently estimates Ψ up to a scale assuming $\Psi_{11} \neq 0$. The latter corresponds to the assumption that the instrument is valid, that is, the instrument and the endogenous variable are correlated.

B Second step: Prediction of ε_t

This section derives the minimum-MSE prediction of the unobserved underlying shocks based on the impact matrix and observed data. In addition, it discusses the relationship to other approaches for shock extraction based on the impact matrix.

B.1 Minimum-MSE prediction of ε_t

Recall that we can obtain the least squares forecast of a state vector in a state-space model using the Kalman filter. Following Hamilton (1994), p. 372, consider the state-space system:

$$\begin{aligned}
\xi_t &= F\xi_{t-1} + v_t \\
y_t &= A'x_t + H'\xi_t + w_t \\
Q &= \mathbb{E}[v_t v_t'] \\
R &= \mathbb{E}[w_t w_t']
\end{aligned}$$

where v_t and w_t are serially and mutually uncorrelated random vectors with zero mean. y_t is a vector of observed variables, ξ_t a vector of unobserved states, and x_t is a vector of exogenous or predetermined variables. F, A', H' are conformable matrices of parameters.

We can use a formula based on the Kalman filter to obtain predictions of these states. We are ultimately interested in $\xi_{t|t} \equiv \mathbb{E}[\xi_t|\Omega_t]$, the linear projection of the unobserved states on the information set $\Omega_t \in \{y_t, y_{t-1}, \ldots, x_t, x_{t-1}, \ldots\}$ on day t. The Kalman filter recursively provides the linear projection in the state-space model (Hamilton, 1994, p. 379):

$$\xi_{t|t} = \xi_{t|t-1} + P_{t|t-1}H(H'P_{t|t-1}H + R)^{-1}(y_t - A'x_t - H'\xi_{t|t-1})$$

where $\xi_{t|t-1}$ is the one-step-ahead prediction of the unobserved states and $P_{t|t-1}$ is the one-step-ahead prediction MSE matrix.

We can simplify the prediction by noting that the unobserved shocks in Model (1) are serially uncorrelated. Therefore, $F = \mathbf{0}$, where $\mathbf{0}$ is a conformable matrix of zeros. It follows that $\xi_{t|t-1} = \mathbf{0}$, the unobserved states are unpredictable based on past data, and $P_{t|t-1} = Q$, the MSE matrix corresponds to the unconditional variance of the shocks in the state equation (see Hamilton, 1994, p. 380). In addition, $H'QH + R = \mathbb{E}[(y_t - A'x_t)(y_t - A'x_t)'] \equiv \Sigma$ is the variance-covariance matrix of the residuals between the data and the predetermined variables. Furthermore, the prediction given data available at time t is equal to the prediction given all data ($\varepsilon_{t|T} = \varepsilon_{t|t}$). This implies that using the Kalman smoother is unnecessary because future observations do not contain useful information about the current unobserved state. It follows that the minimum-MSE prediction given all available information is:

$$\xi_{t|T} = QH\Sigma^{-1}(y_t - A'x_t)$$

This prediction will generally be associated with some uncertainty. We can derive the MSE of the Kalman filter prediction (see Hamilton, 1994, p. 380):

$$P_{t|T} = Q - QH(H'QH + R)^{-1}H'Q$$

Based on these preliminaries, it is straightforward to derive a prediction formula for the unobserved underlying shocks in Model (1).

Proof. Let us cast the model in state-space form for $t \in P$:

$$\begin{aligned}
\xi_t &= \varepsilon_t \\
y_t &= \alpha + \Psi \varepsilon_t + \Gamma v_t + \Phi(L) x_{t-1} \\
y_t - A' x_t &= y_t - \alpha - \Phi(L) x_{t-1} = u_t \\
\Sigma &= \mathbb{E}[u_t u_t'] = \Sigma \\
Q &= \mathbb{E}[\varepsilon_t \varepsilon_t'] = \Sigma_\varepsilon \\
R &= \mathbb{E}[\Gamma v_t v_t' \Gamma'] = \Gamma \Sigma_v \Gamma'
\end{aligned}$$

Therefore, the minimum-MSE prediction of the unobserved shocks is:

$$\varepsilon_{t|T} = \Sigma_\varepsilon \Psi' \Sigma^{-1} u_t \quad \text{for } t \in P$$

with MSE matrix:

$$MSE_{t|T} = \Sigma_\varepsilon - \Sigma_\varepsilon \Psi' \Sigma^{-1} \Psi \Sigma_\varepsilon \quad \text{for } t \in P$$

□

This implies that we can compute the minimum-MSE prediction of the shocks based on the variance-covariance matrix of the structural shocks (Σ_ε), the impact matrix (Ψ), the variance-covariance matrix on days with policy events $\Sigma \equiv \Psi \Sigma_\varepsilon \Psi' + \Gamma \Sigma_v \Gamma'$, and the residuals

(u_t) of the regression of y_t on the information set.

B.2 Scale of predicted shocks

Heteroskedasticity and high-frequency identification schemes estimate Ψ using a unit effect normalization (see, e.g., Stock and Watson, 2018). However, we can compute the minimum-MSE predictions up to a scale based on an estimate of each column of Ψ, which is arbitrarily scaled. Let $\tilde{\Psi} = \Psi S$, where S is an arbitrary diagonal matrix normalizing each column of Ψ. Then we can cast Model (1) in state-space form for $t \in P$:

$$\begin{aligned}
\xi_t &= \tilde{\varepsilon}_t = S^{-1}\varepsilon_t \\
u_t &= \tilde{\Psi}\tilde{\varepsilon}_t + \Gamma v_t \\
\tilde{Q} &= \mathbb{E}[\tilde{\varepsilon}_t\tilde{\varepsilon}_t'] = \Sigma_{\tilde{\varepsilon}} = S^{-1}\Sigma_{\varepsilon}S^{-1\prime} \\
R &= \mathbb{E}[\Gamma v_t v_t'\Gamma'] = \Gamma\Sigma_v\Gamma'
\end{aligned}$$

The model that uses the normalization is observationally equivalent to the true model because $\tilde{\Psi}\tilde{\varepsilon}_t = \Psi S S^{-1}\varepsilon_t = \Psi\varepsilon_t$. The minimum-MSE prediction of the unobserved shocks is:

$$\tilde{\varepsilon}_{t|T} = \Sigma_{\tilde{\varepsilon}}\tilde{\Psi}' \underbrace{\left(\tilde{\Psi}\Sigma_{\tilde{\varepsilon}}\tilde{\Psi}' + \Gamma\Sigma_v\Gamma'\right)^{-1}}_{\Sigma^{-1}} u_t \quad \text{for } t \in P$$

with MSE matrix:

$$\tilde{MSE}_{t|T} = \Sigma_{\tilde{\varepsilon}} - \Sigma_{\tilde{\varepsilon}}\tilde{\Psi}' \underbrace{\left(\tilde{\Psi}\Sigma_{\tilde{\varepsilon}}\tilde{\Psi}' + \Gamma\Sigma_v\Gamma'\right)^{-1}}_{\Sigma^{-1}} \tilde{\Psi}\Sigma_{\tilde{\varepsilon}} \quad \text{for } t \in P$$

This follows from $\tilde{\Psi}\Sigma_{\tilde{\varepsilon}}\tilde{\Psi}' + \Gamma\Sigma_v\Gamma' = \Psi S S^{-1}\Sigma_{\varepsilon} S^{-1'} S'\Psi' + \Gamma\Sigma_v\Gamma' = \Sigma$. Therefore, we can predict unobserved shocks up to a scale based on a scaled estimate of the impact matrix $\tilde{\Psi}$, the variance-covariance matrix of the data on event days Σ, and using an arbitrary normalization for the scaled variance of the policy event shocks $\Sigma_{\tilde{\varepsilon}}$.

B.3 Fama-MacBeth regressions

Bu et al. (2021) propose a two-step approach to estimate a one-dimensional unobserved monetary policy shock. First, they estimate the causal impact of a monetary policy shock on interest rate changes across the term structure using the heteroskedasticity-IV approach by Rigobon and Sack (2004). Second, in the spirit of Fama and MacBeth (1973), they perform a cross-sectional regression of interest rate changes along the term structure on the impact vector for every day with an FOMC decision. The OLS coefficients of these regressions are then proportional to the underlying unobserved monetary policy shocks. Although their approach works within their specific model structure, this section shows that it requires restrictive assumptions and therefore does not yield the optimal prediction of the unobserved shocks if these assumptions are not met.

Suppose that we know each column of Ψ up to a scale, $\tilde{\Psi} = \Psi S$, where S is a conformable diagonal matrix with scaling factors. The OLS estimator of y_t on $\tilde{\Psi}$ on day t reads:

$$\hat{\varepsilon}_t = (\tilde{\Psi}'\tilde{\Psi})^{-1}\tilde{\Psi}'y_t = S^{-1}\varepsilon_t + (\tilde{\Psi}'\tilde{\Psi})^{-1}\tilde{\Psi}'\Gamma v_t$$

The OLS estimate depends proportionally on the true shocks and an error term. OLS provides an unbiased prediction of the true shocks, up to a scale S^{-1}, if $\tilde{\Psi}'\Gamma = \mathbf{0}$. That is, we need an orthogonality condition between the impact matrices of the structural shocks and other shocks, respectively. If this assumption is violated, the shock estimates will suffer from an unobserved variables bias because we fail to control for variation in Γ in the cross-sectional regression.[22] If the orthogonality assumption does not hold, the extracted shocks are contaminated by other shocks occurring on policy event days (v_t).

For illustration, suppose that ε_t and v_t are one-dimensional shocks. Then $\tilde{\Psi}$ and Γ are impact vectors. Then we have:

$$\hat{\varepsilon}_t = \frac{1}{s}\varepsilon_t + \frac{\sum_{i=1}^{N}\tilde{\psi}_i\gamma_i}{\sum_{i=1}^{N}\tilde{\psi}_i^2}v_t$$

where $\tilde{\psi}_i$ and γ_i are the i^{th} element of the impact vectors. This implies that we can recover ε_t up to a scale only if $\gamma_i = 0$ for all i, or $\sum_{i=1}^{N}\tilde{\psi}_i\gamma_i = 0$. The former implies that no other shocks occur on policy event days, which is a very restrictive assumption. The latter implies that the impact vectors are orthogonal, which may be restrictive depending on the specific application.

---

[22] In monetary policy applications, for example, the assumption is violated if long-term interest rates respond less strongly than short-term interest rates to both monetary policy and productivity shocks. Whether this assumption is violated in practice is an empirical question, and we will compare the resulting shock series with both approaches to existing multi-dimensional high-frequency surprises.

C Data

Tables 8 and 9 summarize the sources for time series, policy events, and control days.

Tab. 8: Time series

Category	**Source**	**Variants**	**Time stamp**	**Comments**
Treasury bill yields	Board of Governors	3M, 6M, 1Y, 2Y, 3Y, 5Y, 7Y, 10Y, 30Y	4pm EST	www.federalreserve.gov/releases/h15/. We do not use the 1M and 20Y yields as they comprise missing data
Federal Funds Rate	Board of Governors		4pm EST	www.federalreserve.gov/releases/h15/
Nominal effective exchange rate	Board of Governors		Noon EST	www.federalreserve.gov/releases/h10/. We linked the discontinued series with FRED identifier DTWEXM with DTWEXAFEGS. Also, we define the NEER as one unit of foreign currency in terms of USD. Therefore, a decline in the exchange rate is an appreciation of the USD.
Stock prices	S&P Dow Jones, NASDAQ OMX	S&P 500, NASDAQ	4pm EST	de.tradingview.com/symbols/SPX/. FRED variable keys: NASDAQCOM, NASDAQ100
VIX	CBOE		Close	FRED variable keys: VIXCLS, VXOCLS
Corporate bond spreads	Moody's	AAA, BAA	Unclear	FRED variable keys: DAAA, DBAAA. We computed the spreads as the difference to the 10Y government bond yield
Corporate bond prices	BofA Merrill Lynch	Total, 1Y-3Y, 3Y-5Y, 10Y-15Y, 15Y+	Close	fred.stlouisfed.org/categories/32413. Mostly start in 1988
Treasury yield skewness	Bauer and Chernov (2024)		End-of-day	The underlying data are treasury futures and options from CME group. Starts in 1988
Industrial production	Board of Governors			FRED variable key: INDPRO
CPI	BLS			FRED variable key: CPIAUCSL

To capture the information across a wider variety of financial market variables, and potentially reduce some noise in the individual series, we compute averages over multiple series.

The short-term rate averages the FFR, as well as the 3M and 6M treasury bill rates. The medium-term interest rate averages the 2Y, 3Y and 5Y treasury bond yields. The long-term rate averages 7Y, 10Y and 30Y treasury bond yields.[23] The stock price index averages the S&P 500, as well as the NASDAQ Composite Index and the NASDAQ 100 Index. The former receives a weight of 0.5, and the other two each receive a weight of 0.25. The corporate bond spread averages the series for AAA and BAA rated bonds. Note that all results remain robust using various versions of the individual series.

In the daily model, we use the announcement dates and monetary policy surprises of Bauer and Swanson (2022). In a robustness test, we use various alternative event days over an extended sample from 1986–2022, and control for various contaminating events (e.g., data releases or decisions by other central banks), based on our own data collection (see Table 9.

In the monthly model, we aggregate the daily financial market variables using the average of daily observations. In addition, we use standard data on seasonally adjusted consumer prices and industrial production.

---

[23] We also estimated models based on data from Gürkaynak et al. (2007) featuring more maturities. The results are qualitatively similar. However, we prefer yields published by the Board of Governors because we know the exact time stamp of the daily data.

Tab. 9: Events

Category	Source	Start	Comments
FOMC announcements	1982 to 1987: www.federalreserve.gov/monetarypolicy/fomc_historical_year.htm, 1988 to 2019: Bauer and Swanson (2022), from 2020: www.federalreserve.gov/monetarypolicy/fomccalendars.htm	1982	
Discount rate changes	Board of Governors	1934	FRED variable key: DISCOUNT
Federal Funds target changes	Board of Governors	1982	FRED variable key. DFEDTAR
Speeches and Congressional Testimony	Up to 1996: alfred.stlouisfed.org/, from 1997: www.federalreserve.gov/newsevents/speeches.htm	1982	
FOMC minutes	www.federalreserve.gov/monetarypolicy/fomccalendars.htm	1988	
ECB decisions	Altavilla et al. (2019)	1999	
BoE decisions	Braun et al. (2024)	1997	
CPI releases	www.bls.gov/bls/news-release/cpi.htm, alfred.stlouisfed.org/	1982	
PPI releases	www.bls.gov/bls/news-release/ppi.htm	1994	
Employment situation releases	www.bls.gov/bls/news-release/empsit.htm	1994	
Employment cost releases	www.bls.gov/bls/news-release/eci.htm	1982	
GDP releases	www.bea.gov/index.php/news/archive?field_related_product_target_id=All&created_1=All&title=gross%20domestic%20product&page=0	1996	Includes first, second and third releases
Industrial production releases	www.federalreserve.gov/releases/g17/release_dates.htm	1982	

D Robustness tests

Tab. 10: Correlations of two-step predictions based on alternative model specifications

	Baseline	$L = 0$	$L = 2$	$L = 4$	Long sample	Excluding other events
Baseline	1.00					
$L = 0$	1.00	1.00				
$L = 2$	1.00	1.00	1.00			
$L = 4$	1.00	1.00	1.00	1.00		
Long sample	0.94	0.94	0.94	0.94	1.00	
Excluding other events	0.99	0.98	0.98	0.98	0.96	1.00

Notes: The table shows correlations of the heteroskedasticity-based two-step predictions using various model specifications. We vary the lag length of the control variables (L), estimate the model on a longer sample with own events, and exclude contaminating events, such as speeches, releases of minutes, data releases, and decisions by other central banks, from the control data set.

Fig. 4: Correlation of two-step predictions with true shocks (conditional heteroskedasticity)

(a) Variance of event shocks **(b)** Accuracy of proxy

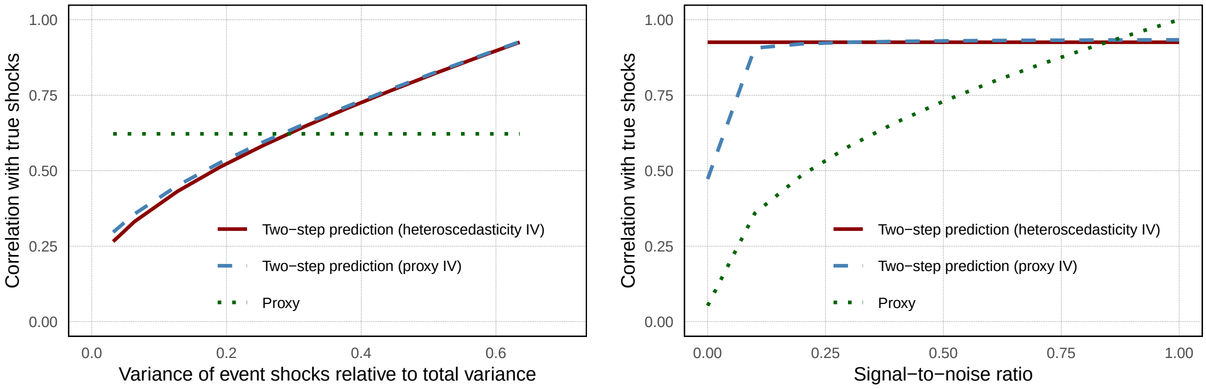

(c) Accuracy of estimates **(d)** Frequency of events

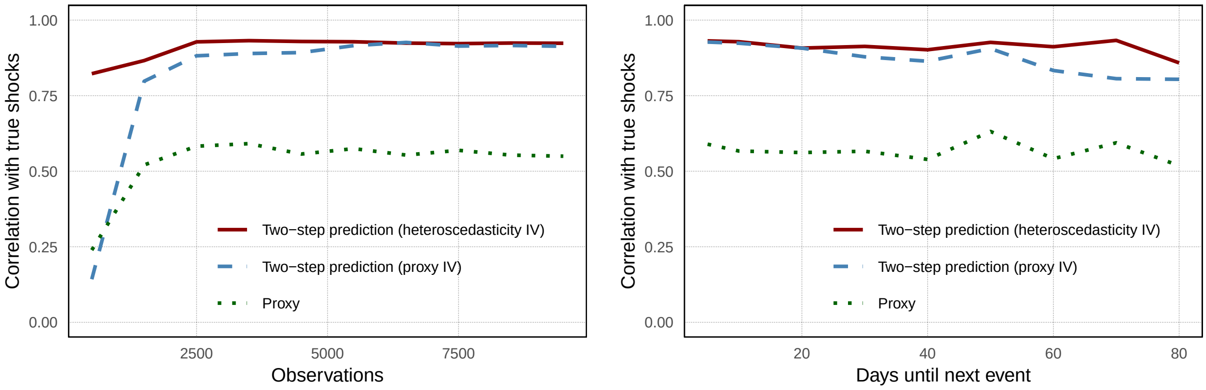

(e) Number of predictors

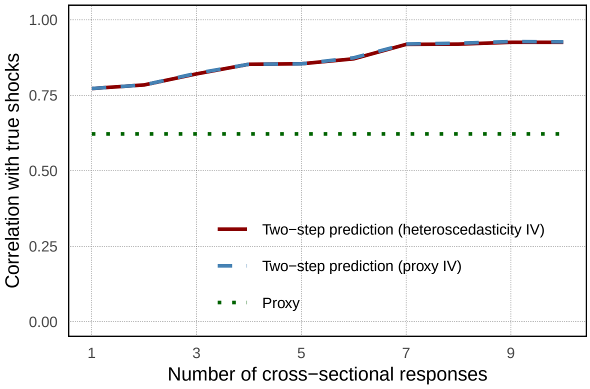

Notes: Correlations of monetary policy shock measures with the true underlying shocks based on simulated data. The structural shocks are simulated with conditional heteroskedasticity, that is, GARCH(1,1) processes with $\sigma_t^2 = \mu + \alpha\zeta_{t-1}^2 + \beta\sigma_{t-1}^2$ and $\mu = 0.1, \alpha = 0.1, \beta = 0.8$, where ζ_t corresponds to the respective policy event (ε_t) or other shock (v_t). Panel a) varies the variance of event shocks relative to the total variance in the simulated medium-term interest rate. Panel b) varies the signal-to-noise ratio of the proxy of the event policy shocks. Panel c) varies the number of observations to estimate the model parameters. Panel d) varies the frequency of events. Panel e) varies the number of impact responses used to predict the unobserved policy event shocks.

Fig. 5: Correlation of two-step predictions with true shocks (fat tails)

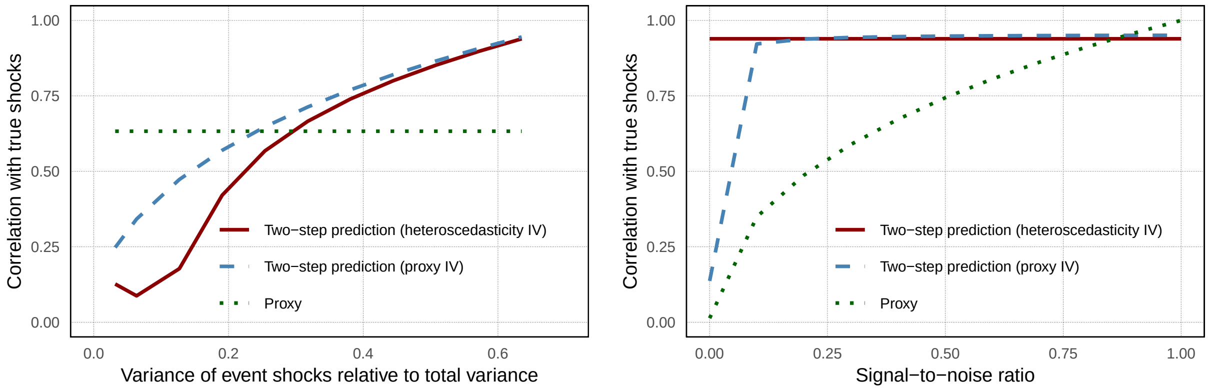

(a) Variance of event shocks

(b) Accuracy of proxy

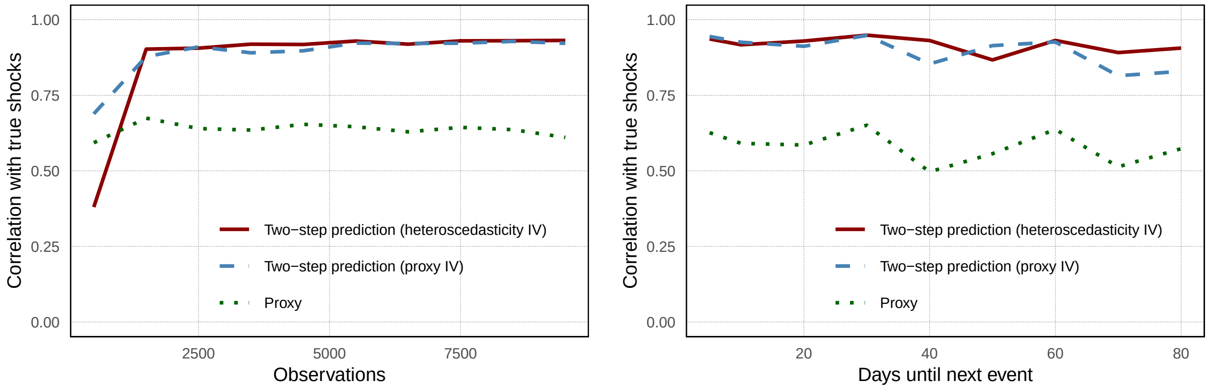

(c) Accuracy of estimates

(d) Frequency of events

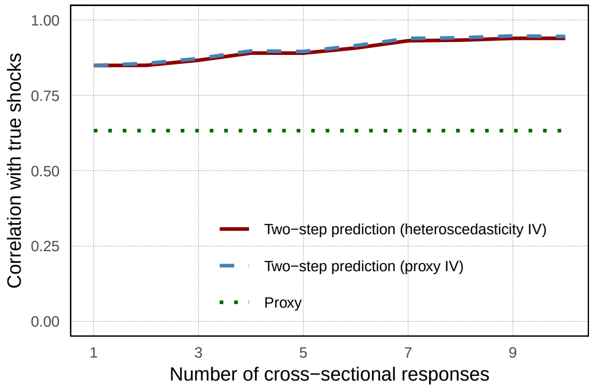

(e) Number of predictors

Notes: Correlations of monetary policy shock measures with the true underlying shocks based on simulated data. The structural shocks are simulated from a t-distribution with 3 degrees of freedom. Panel a) varies the variance of event shocks relative to the total variance in the simulated medium-term interest rate. Panel b) varies the signal-to-noise ratio of the proxy of the event policy shocks. Panel c) varies the number of observations to estimate the model parameters. Panel d) varies the frequency of events. Panel e) varies the number of impact responses used to predict the unobserved policy event shocks.