Economics Letters 268 (2026) 113091

Contents lists available at ScienceDirect

Economics Letters

journal homepage: www.elsevier.com/locate/ecolet

Multiple monetary policy shocks from daily data: A heteroskedasticity IV approach ☆

Marc Burri [a], Daniel Kaufmann [a,b,*]

[a] University of Neuchâtel, Institute of Economic Research, Rue A.-L. Breguet 2, CH-2000, Neuchâtel, Switzerland
[b] KOF Swiss Economic Institute, ETH Zurich, Switzerland

ARTICLE INFO

Dataset link: Replication files (Original data)

JEL classification:
C3
E3
E4
E5

Keywords:
Monetary policy shocks
Causal effects
Forward guidance
Heteroskedasticity
High-frequency
Instrumental variables

ABSTRACT

We extend the heteroskedasticity IV estimator of Rigobon and Sack (2004) to multiple monetary policy shocks by imposing recursive zero restrictions on the impact matrix. Unlike high-frequency identification, the approach requires neither intraday tick data nor precise announcement timestamps, making it applicable to countries and historical periods where such data are unavailable. Applied to US FOMC announcements, we find causal effects similar to those of high-frequency identification. The heteroskedasticity-based instrument passes weak-instrument tests for the target shock, whereas results are mixed for high-frequency surprises. For the path shock, we obtain strong instruments for the majority of specifications and the underlying shock series is highly correlated with the high-frequency counterpart.

1. Introduction

We propose a heteroskedasticity-based instrumental variables (HET-IV) approach to identify the causal effects of monetary policy target and path shocks from daily financial market data.[1] The approach extends the IV estimator of Rigobon and Sack (2004) from one to multiple shocks by imposing recursive zero restrictions on the impact matrix along the term structure of interest rates, which mirror those used for constructing multi-dimensional high-frequency instruments (HF-IV) (see, e.g., Gürkaynak et al., 2005; Swanson, 2021). Unlike HF-IV, the approach does not require intraday tick data or knowledge of the exact timing of announcements, making it applicable to countries and historical periods where such information is unavailable. The recursive IV structure yields a closed-form estimator to which the weak-instrument tests of Lewis and Mertens (2025) apply directly, enabling straightforward instrument-strength evaluation for multi-dimensional shocks identified through HET-IV.

Related work identifies multiple monetary policy shocks using heteroskedasticity (Wright, 2012; Canetg and Kaufmann, 2022), non-Gaussianity (Jarociński, 2024), multi-dimensional HF-IVs (Swanson, 2021; Altavilla et al., 2019), or combining external instruments, heteroskedasticity, and parameter restrictions (Schlaak et al., 2023; Carriero et al., 2024; Jarocinski and Karadi, 2025).[2] We add to this literature by proposing a closed-form IV estimator, exploiting the change in the variance–covariance matrix of daily financial variables combined with recursive zero restrictions to disentangle multiple monetary policy shock dimensions. Although high-frequency data sets are becoming increasingly available for various countries (Altavilla et al., 2019; Acosta et al., 2025; Braun et al., 2024; Emeksiz et al., 2026), data

☆ We thank Gianluca Benigno, Kenza Benhima, Aurélien Eyquem, Jean-Marie Grether, Kevin Kotzé, Leif Anders Thorsrud, Jean-Paul Renne, Eric Swanson, 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, and the University of Lausanne 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 "Causal effects of multiple monetary policy dimensions: A heteroskedasticity IV approach".

* Corresponding author at: University of Neuchâtel, Institute of Economic Research, Rue A.-L. Breguet 2, CH-2000, Neuchâtel, Switzerland.
E-mail addresses: marc.burri91@gmail.com (M. Burri), daniel.kaufmann@unine.ch (D. Kaufmann).

[1] The estimator is implemented in the R package hetiv which can be downloaded from https://dankaufmann.github.io/hetiv.

[2] Another strand of the literature uses heteroskedasticity to obtain overidentified models, such that the validity of traditional identifying assumptions can be tested (see, e.g., Rigobon and Sack, 2004; Lütkepohl and Schlaak, 2022; Schlaak et al., 2023).

https://doi.org/10.1016/j.econlet.2026.113091
Received 2 April 2026; Received in revised form 8 June 2026; Accepted 8 June 2026
Available online 2 July 2026

are often only available over relatively short sample periods. Because daily data are available over longer sample periods, our approach is especially useful when analyzing causal effects of monetary policy or other announcements through their financial market effects over longer historical episodes.

Since the seminal contribution by Sims (1980), recursive zero restrictions have been used to identify monetary policy shocks in structural vector autoregressions (see, e.g., Christiano et al., 1999; Stock and Watson, 2001). Our approach resembles this strand of the literature by imposing zero restrictions on the impact matrix. However, zero restrictions are only required to disentangle multiple dimensions of monetary policy shocks, while the overall effect of a linear combination of those shocks is identified via heteroskedasticity. Recursive zero restrictions are plausible when there is a clear hierarchy of how multiple shocks are transmitted to the variables on impact—here, target shocks move all variables while a path shock leaves the short-term interest rate unchanged within the same working day.[3]

Applying the approach to US FOMC announcements, we find causal effects similar to those based on HF-IVs by Swanson (2021). A key advantage of HET-IV is its strong instrument performance: for the target shock, the heteroskedasticity- and autocorrelation-robust (HAR) test statistic of Montiel Olea and Pflueger (2013) ranges from 48.0 to 148.2 in various specifications, well above the critical value of 14.2, whereas the HF-IV fails the same test in two out of four specifications. For the path shock, the HET-IVs pass the weak-instrument tests in most specifications, and the results are close in others. Meanwhile, the HF-IVs pass all weak-instrument tests for the path shock. In addition, the point estimates of the responses to a path shock are close, and we generally cannot reject the null hypothesis of equality of the responses based on HET-IV and HF-IV. This suggests that potential weak-instrument biases in the HET-IV responses are small. While Lewis (2022) documents that HET-IVs based on high-frequency data yield stronger instruments than daily data, our results suggest that daily data can still yield strong multi-dimensional instruments on longer sample periods.

In what follows, Section 2 presents the model. Section 3 reports results for US FOMC announcements. Section 4 concludes.

2. Identification of multiple shocks with HET-IV

2.1. Model

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

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

where α is a vector of constants, ε_t is a vector of E structural shocks on policy event days (P), and υ_t is a vector of R other shocks on policy event as well as control days (P and C). 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, 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.

2.2. Estimation and identification

It is well known that, if the monetary policy shock is one-dimensional ($E = 1$), we can estimate Ψ up to a scale using the following instrumental variable (Rigobon and Sack, 2004; Lewis, 2022)[4]:

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

where, without loss of generality, we treat variable y_{1t} as the endogenous variable, u_{1t} is the residual of y_{1t} with $u_t = y_t - \Phi(L)x_{t-1}$, $\mathbf{1}(t \in P)$ and $\mathbf{1}(t \in C)$ denote indicator functions that equal one on policy event and control days, respectively, and zero otherwise, and T, T_P, and T_C are the number of total, policy event and control days.[5]

However, for $E > 1$, the standard IV fails. Let Z_{et} be the heteroskedasticity-based instrument constructed using the residual for variable e in Model (1). In addition, let Ψ_{ie} denote the ith row and eth column of Ψ. Finally, let u_i and Z_e denote the column vectors of all observations of the residual i and instrument e, respectively.

Proposition 1. *For $E > 1$ and $e = 1 \ldots E$*

$$\Psi^{IV}_{ie} = [Z'_e u_e]^{-1} Z'_e u_i$$

is an inconsistent estimator of Ψ_{ie}.

Proof. See Online Appendix B.3. □

Intuitively, using only one IV, independent of which variable we use to construct it, is inconsistent because multiple shocks affect the endogenous variables on event days and the instrument is correlated with all of those shocks.

However, using additional identifying assumptions, a modified HET-IV consistently estimates each column of Ψ up to a scale.

Assumption 1. The leading $E \times E$ principal submatrix of Ψ is lower-triangular.

Proposition 2. *Under Assumption 1, the last element of*

$$\Psi^{IV}_{ie} = [(Z_1 \ \ldots \ Z_e)'(u_1 \ \ldots \ u_e)]^{-1}(Z_1 \ \ldots \ Z_e)'u_i$$

is a consistent estimator of Ψ_{ie} up to a scale.

Proof. See Online Appendix B.4. □

This implies that we can recursively estimate each column of Ψ up to a scale. To estimate the column e, we use the first e instruments for the first e endogenous variables. Intuitively, the second instrument is correlated with both the first and second structural shocks. Therefore, only using the second instrument does not allow us to disentangle the two dimensions. By including the first instrument, which is by assumption only correlated with the first shock, we control for the variation in the second variable that is caused by the first shock.

3. Multi-dimensional monetary policy shocks in the US

We apply our approach to US data because we can compare the results to existing multi-dimensional high-frequency surprises.

3.1. Identifying assumptions

The order of the variables in y_t is key for identification. We impose recursive zero restrictions along the term structure of interest rates: a target shock affects all financial market variables on impact, whereas a path shock does not affect the short-term interest rate contemporaneously. The restriction is economically plausible: a path shock should

[3] Recursive zero restrictions are inappropriate in applications where the shocks have contemporaneous effects across all variables, in which case alternative restrictions (e.g., sign restrictions; Uhlig, 2005) may provide a plausible alternative.

[4] See also Online Appendix B.2.

[5] We follow Lewis (2022) and construct the HET-IV from consistently estimated residuals. Therefore, the standard errors of the IV estimator are subject to a generated-regressor problem. The magnitude of the problem decreases with sample size. In our application with daily data, due to the large sample size this problem is negligible.

move only forward-looking variables on the same working day, since the central bank's target rate is fixed until the next meeting. The restriction also mirrors the factor rotation approach used in standard HF-IV schemes (Gürkaynak et al., 2005; Swanson, 2021), giving the shocks a similar economic interpretation as their HF-IV counterparts.[6]

3.2. Data

In y_t we include a short- and medium-term interest rate, as well as a corporate bond spread, an effective exchange rate, stock prices, and the VIX.[7] As control variables (x_{t-1}) we include one lag of the same variables and the implied treasury yield skewness by Bauer and Chernov (2024).[8] All interest rates, spreads and the VIX are included in first-differences. The treasury yield skewness is included in levels. All other variables are included in log-differences multiplied by 100. The exchange rate is defined as one unit of foreign currency measured in USD (a decline is an appreciation of the Dollar). As policy events, we use FOMC announcement dates by Bauer and Swanson (2022) from 1988–1993 and by Acosta et al. (2025) from 1994–2025. All other weekdays, excluding holidays, are used as control days.

3.3. Dynamic causal effects

We estimate dynamic causal effects in a long-difference local projection framework (Jordà, 2005).[9] Specifically, for each horizon $h = 0, 1, \ldots, H$, outcome variable $i = 1, \ldots, N$, and shock dimension $e = 1, \ldots, E$ we estimate:

$$y_{i,t+h} - y_{i,t-1} = \alpha_{ieh} + \sum_{j=1}^{e} \psi_{ijh} \left(i_{j,t} - i_{j,t-1}\right) + \Phi_{ieh} x_{t-1} + u_{ie,t+h}\,, \tag{2}$$

where i_{et} is the interest rate corresponding to shock e, treated as endogenous and instrumented by Z_{et} from Proposition 2. That is, for $e = 1$ we include the short-term interest rate instrumented by Z_{1t} and for $e = 2$ we include the short-term and medium-term interest rate instrumented by Z_{1t} and Z_{2t}. The impulse response of variable i to shock e at horizon h is $\hat{\psi}_{ieh}$, where we normalize the estimate to a 25 basis point increase in the corresponding interest rate at $h = 0$. Inference is based on HC standard errors (Montiel Olea et al., 2025).

Fig. 1 panel (a) shows the causal effect of a target shock, for which the instrument is constructed based on the short-term interest rate. A 25 bp increase in the short-term interest rate, caused by a target shock, leads to a smaller increase in medium-term interest rates, a decline in the corporate bond spread and an appreciation of the US Dollar. Meanwhile, stock prices and the VIX do not show a statistically significant response.

The path shock shown in panel (b) has qualitatively similar effects on most variables. However, there are two important differences. First, the short-term interest rate does not change, as assumed, on impact. Meanwhile, the 25 bp increase in the medium-term interest rate leads to a stronger response of all other variables than the target shock. In particular, the US Dollar appreciates more and stock prices fall more. In addition, the VIX shows a significant increase for some single horizons.

[6] Altavilla et al. (2019) additionally differentiate between press releases, which may comprise a target shock, and press conferences, which may comprise a path shock.

[7] The data sources are given in Online Appendix A.

[8] We include this variable because (Bauer and Swanson, 2022) have shown it to predict high-frequency monetary policy surprises.

[9] Compared to a level specification, this framework is more robust when estimating persistent responses (see Piger and Stockwell, 2025). Compared to vector autoregressions, local projections are 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 of local projections and VARs.

3.4. Comparison with high-frequency identification and weak-instrument tests

We estimate dynamic causal effects using FFR and forward guidance surprises by Swanson (2021) as instruments and compare them to heteroskedasticity-based responses (see Fig. 2).[10]

For the path shock (panel b), we find qualitatively identical point estimates. We also bootstrap p-values for a test of equality of the impulse responses identified via HF-IV and HET-IV.[11] Across the vast majority of horizons, we do not reject the null hypothesis of equality of the responses. Interestingly, the impact effect of the HF-IV path shock on the short-term interest rate is not statistically significantly different from zero. This lends additional credibility to our recursive identifying assumption when using HET-IV.

Differences are more pronounced for the target shock, where responses using the HF-IV are larger in absolute value for some variables. Again, we formally test equality of the impulse responses. Even though the HF-IV confidence intervals are wider, we reject the null of equality somewhat more often than for the path shock. We therefore next turn to the question of why the differences are more pronounced for the target shock.

We present weak-instrument tests, which additionally explain the differences across approaches. The tests follow Montiel Olea and Pflueger (2013), Lewis (2022), and Lewis and Mertens (2025) and account for heteroskedasticity and autocorrelation.[12]

Table 1 reports HAR weak-instrument tests for the target shock. Because results may differ depending on the choice of the endogenous variable, we present various specifications. Panel (a) shows that the heteroskedasticity-based instrument passes the test comfortably for all short-term interest rate specifications, with test statistics of 48.0–148.2 against a critical value of 14.2. Panel (b) reveals that high-frequency FFR surprises fail the same test in two out of four specifications. In particular, the specification with the average short-term rate suffers from a weak-instrument problem, which explains the wider confidence intervals seen in Fig. 2 and differences of the point estimates.[13]

Table 2 reports tests for the joint two-dimensional case based on Lewis and Mertens (2025). This is a direct generalization of Montiel Olea and Pflueger (2013) to multiple endogenous regressors and instruments. The strength of the HET-IV for the second dimension depends on the choice of the HET-IV for the first dimension. We therefore report multiple combinations. Panel (a) shows that the HET-IVs jointly pass in most specifications. In particular, using the average short-term interest rate for the first dimension yields consistently strong instruments for the second dimension. These results are relevant for applied researchers, who aim to obtain strong instruments for multiple monetary policy shocks: the target shock is robustly identified, while

[10] We do so by estimating an equation similar to Eq. (2). There are two differences. First, we use the HF-IVs. Second, we do not impose the recursive zero restriction. Online Appendix C shows that the results remain virtually identical when imposing the zero restriction.

[11] See Online Appendix C.

[12] For the path shock, they also give indirect evidence of whether two dimensions exist. If no path shock exists the first instrument already captures the relevant variation, and adding a second instrument provides no additional information.

[13] Lewis (2022) shows that in an application using data from Nakamura and Steinsson (2018) daily data yield weaker heteroskedasticity-based instruments than intraday data. While we do not dispute this result, panel (a) of Table 1 demonstrates that heteroskedasticity-based instruments can be strong even at daily frequency. However, the results are not directly comparable because we use a longer sample, examine multiple monetary policy shock dimensions, and control for lags of dependent variables. The latter may be especially important as daily financial market data and high-frequency surprises may be correlated with information available on FOMC announcement days (Bauer and Swanson, 2023, 2022).

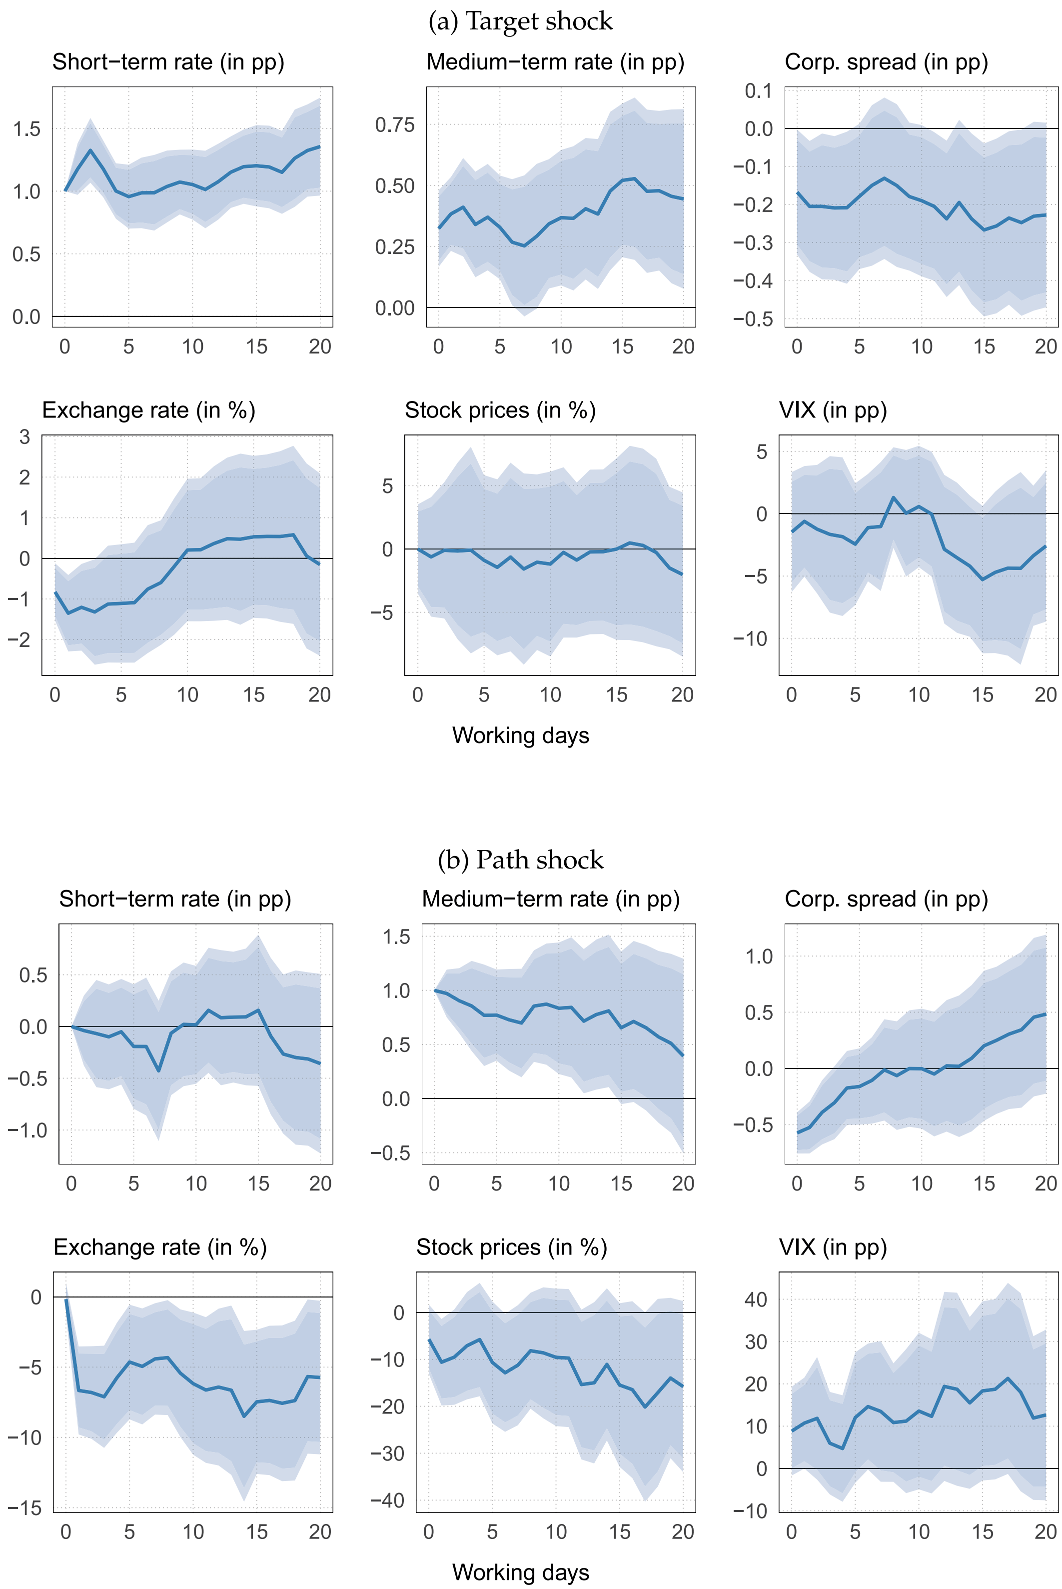

Fig. 1. Dynamic causal effects using HET-IVs.

Notes: Impulse responses to a target (panel a) and path (panel b) monetary policy shock identified via HET-IV. The responses are normalized to a 25 bp increase in the short-term and medium-term interest rate, respectively. The horizontal axis measures working days (excluding weekends and holidays). The model is 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 shaded areas show 90% and 95% confidence intervals, which are based on HC standard errors.

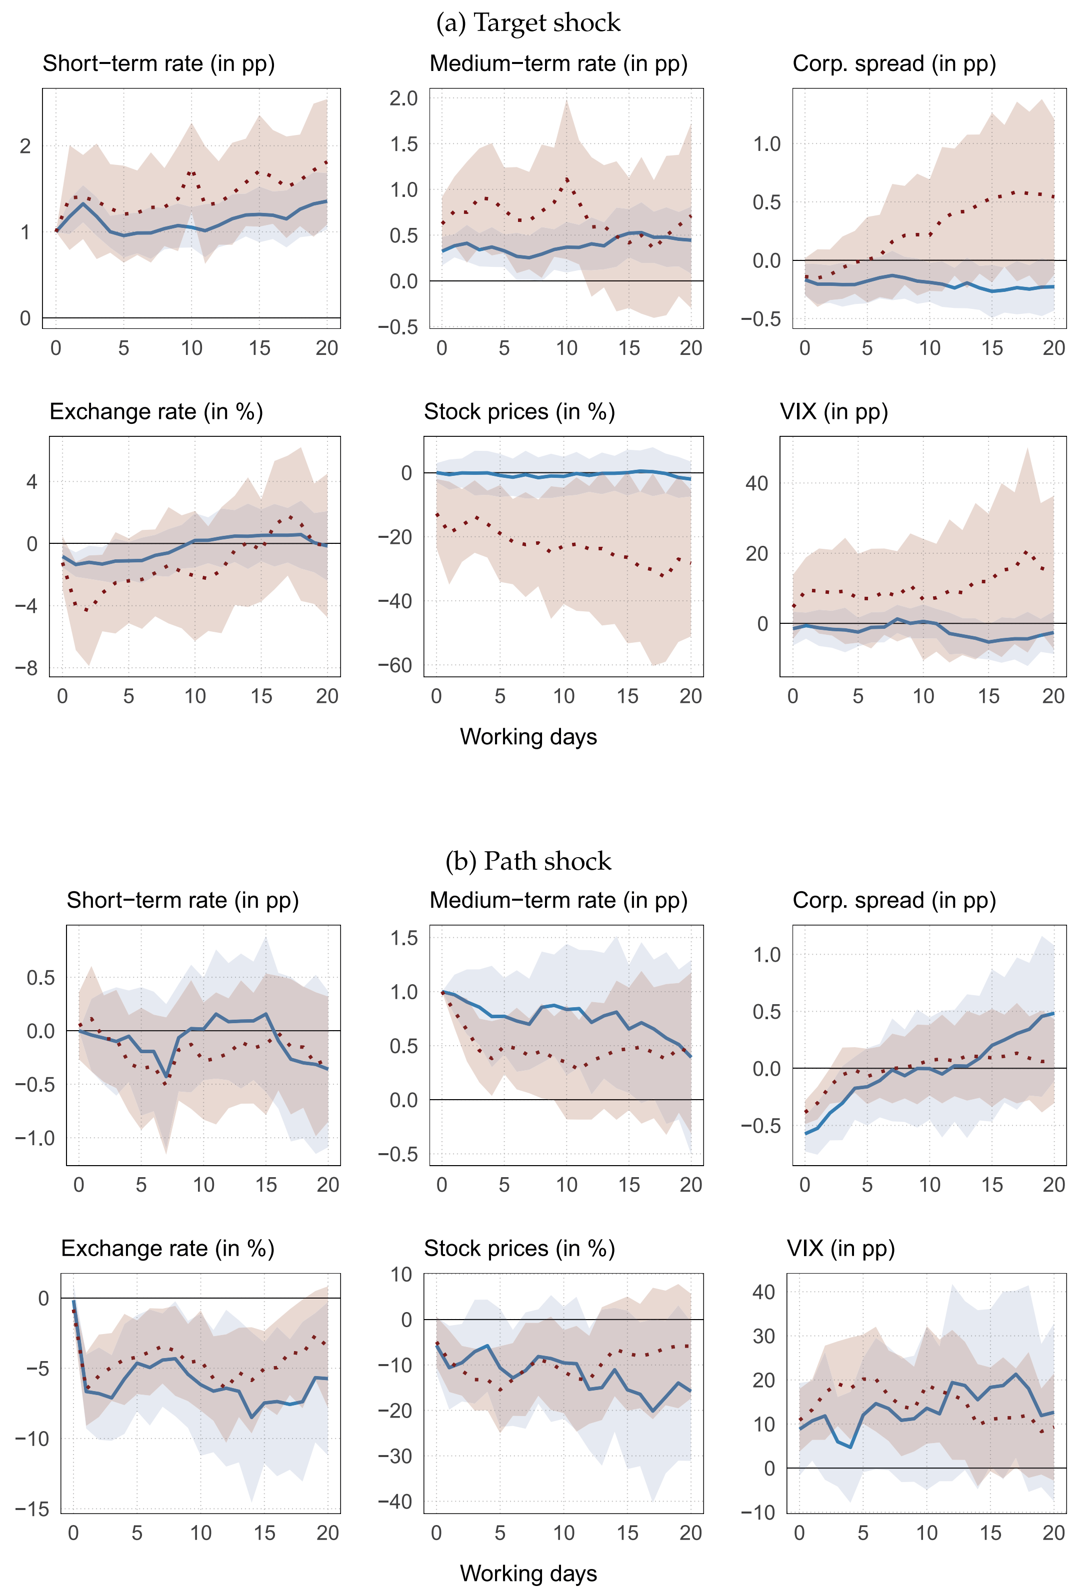

Fig. 2. Comparison of HET-IVs with HF-IVs.
Notes: Impulse responses to a target (panel a) and path (panel b) monetary policy shock identified via HF-IV (red dotted line) and HET-IV (blue solid line). Shaded areas show 90% confidence intervals (blue shading for HET-IV, red for HF-IV), based on HC standard errors. See also notes to Fig. 1.

Table 1
Weak-instrument tests for target shocks.

(a) HET-IV				
	FFR	3M	6M	Average
Test statistic	59.5	48.0	90.6	148.2
Critical value	14.2	14.2	14.2	14.2
(b) HF-IV				
	FFR	3M	6M	Average
Test statistic	5.6	16.3	27.6	12.9
Critical value	14.2	14.2	14.2	14.2

Notes: Weak-instrument tests robust to heteroskedasticity and autocorrelation (Montiel Olea and Pflueger, 2013; Lewis, 2022; Lewis and Mertens, 2025). Each column varies the endogenous variable. For heteroskedasticity-based identification, we also vary the variable to construct the instrument. For high-frequency identification, we always use the FFR surprises by Swanson (2021). The significance level is set to 5% and the tolerance level to 10%.

Table 2
Weak-instrument tests for path shocks.

(a) HET-IV				
Target shock	Path shock			
	2Y	3Y	5Y	Average
FFR	29.0	30.4	32.9	31.3
Critical value	30.5	30.3	29.3	30.2
3M	34.9	34.2	31.5	36.1
Critical value	31.3	32.3	33.2	32.3
6M	35.2	34.8	28.6	35.6
Critical value	31.0	31.7	32.7	31.6
Average	38.5	36.5	35.0	38.2
Critical value	30.6	31.1	30.6	30.8
(b) HF-IV				
	2Y	3Y	5Y	Average
Test statistic	110.5	111.6	76.9	111.6
Critical value	14.2	14.2	14.2	14.2

Notes: Weak-instrument tests robust to heteroskedasticity and autocorrelation (Montiel Olea and Pflueger, 2013; Lewis, 2022; Lewis and Mertens, 2025). Panel (a) uses the test for multiple instruments and multiple endogenous regressors by Lewis and Mertens (2025). Each row varies the endogenous variable and instrument for the first dimension (target shock). Each column varies the endogenous variable and instrument for the second dimension (path shock). Panel (b) uses the test for one instrument and one endogenous regressor by Montiel Olea and Pflueger (2013). For high-frequency identification, we always use the forward guidance surprises by Swanson (2021). The significance level is set to 5% and the tolerance level to 10%.

path-shock inference is slightly more sensitive to the choice of the endogenous variable. For completeness, Panel (b) reports that the high-frequency forward guidance surprises easily pass the one-dimensional weak-instrument test.[14]

To further examine whether the two approaches identify similar underlying shocks, we predict the shocks using the approach by Burri and Kaufmann (2026).[15] For HET-IV, the target and path shocks are virtually uncorrelated (0.07), as we would expect. In addition, the HET-IV path shock is highly correlated with the HF-IV path shock (0.90), confirming that both approaches identify similar underlying shocks. The HET-IV and HF-IV target shocks exhibit a lower correlation (0.64), reflecting the weak-instrument problem for the HF-IV.

4. Conclusions

We propose a heteroskedasticity IV approach to identify the causal effects of multiple monetary policy shocks from daily financial market data. The closed-form recursive estimator extends (Rigobon and Sack, 2004) to multiple dimensions and enables direct weak-instrument testing via (Lewis and Mertens, 2025). Applied to US FOMC announcements, the approach yields causal effects similar to high-frequency identification while outperforming it on instrument-strength tests for the target shock. For the path shock we obtain similar point estimates for the impulse responses, and similar shock predictions. Key extensions—separately identifying information effects (Nakamura and Steinsson, 2018) and measuring impacts on low-frequency macroeconomic variables via the predicted shock series of Burri and Kaufmann (2026)—are left for future research.

[14] Testing the strength of the instruments jointly yields mixed results.
[15] The results are shown in Online Appendix C.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the authors used Claude, Writefull, and Copilot in order to obtain feedback on structure and grammar, and support with coding in R. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.econlet.2026.113091.

Data availability

Data and replication files are available on the corresponding author's Github Repository

Replication files (Original data) (Github)

References

Acosta, M., Ajello, A., Bauer, M., Loria, F., Miranda-Agrippino, S., 2025. Financial market effects of FOMC communication: Evidence from a new event-study database. Working Paper 2025–30, Federal Reserve Bank of San Francisco.

Altavilla, C., Brugnolini, L., Gürkaynak, R.S., Motto, R., Ragusa, G., 2019. Measuring euro area monetary policy. J. Monet. Econ. 108, 162–179. http://dx.doi.org/10.1016/j.jmoneco.2019.08.016.

Bauer, M., Chernov, M., 2024. Interest rate skewness and biased beliefs. J. Financ. 79 (1), 173–217. http://dx.doi.org/10.1111/jofi.13276.

Bauer, M.D., Swanson, E.T., 2022. A reassessment of monetary policy surprises and high-frequency identification. In: Eichenbaum, M., Hurst, E., Ramey, V.A. (Eds.), NBER Macroeconomics Annual 2022, Volume 37. University of Chicago Press, National Bureau of Economic Research.

Bauer, M.D., Swanson, E.T., 2023. An alternative explanation for the "Fed information effect". Am. Econ. Rev. 113 (3), 664–700. http://dx.doi.org/10.1257/aer.20201220.

Braun, R., Miranda-Agrippino, S., Saha, T., 2024. Measuring monetary policy in the UK: The UK monetary policy event-study database. J. Monet. Econ. http://dx.doi.org/10.1016/j.jmoneco.2024.103645.

Burri, M., Kaufmann, D., 2026. Measuring monetary policy shocks. Working Paper 24–03, IRENE Institute of Economic Research, University of Neuchâtel.

Canetg, F., Kaufmann, D., 2022. Overnight rate and signalling effects of central bank bills. Eur. Econ. Rev. 143, 104060. http://dx.doi.org/10.1016/j.euroecorev.2022.104060.

Carriero, A., Marcellino, M., Tornese, T., 2024. Blended identification in structural VARs. J. Monet. Econ. 146, 103610. http://dx.doi.org/10.1016/j.jmoneco.2024.103610.

Christiano, L.J., Eichenbaum, M., Evans, C.L., 1999. Monetary policy shocks: What have we learned and to what end? In: Handbook of Macroeconomics, vol. 1, Elsevier, pp. 65–148. http://dx.doi.org/10.1016/S1574-0048(99)01005-8.

Emeksiz, E.O., Kamber, G., Otten, J., Pasricha, G.K., 2026. Still packing a punch: Monetary policy transmission in a new cross-country high-frequency dataset. IMF Working Paper 2026/018, International Monetary Fund.

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

Jarociński, M., 2024. Estimating the Fed's unconventional policy shocks. J. Monet. Econ. 144, 103548. http://dx.doi.org/10.1016/j.jmoneco.2024.01.001.

Jarocinski, M., Karadi, P., 2025. Disentangling monetary policy, central bank information, and Fed response to news shocks. Discussion Paper DP19923, CEPR.

Jordà, Ò., 2005. Estimation and inference of impulse responses by local projections. Am. Econ. Rev. 95 (1), 161–182. http://dx.doi.org/10.1257/0002828053828518.

Lewis, D.J., 2022. Robust inference in models identified via heteroskedasticity. Rev. Econ. Stat. 104 (3), 510–524. http://dx.doi.org/10.1162/rest_a_00963.

Lewis, D.J., Mertens, K., 2025. A robust test for weak instruments for 2SLS with multiple endogenous regressors. Rev. Econ. Stud. http://dx.doi.org/10.1093/restud/rdaf103.

Lütkepohl, H., Schlaak, T., 2022. Heteroscedastic proxy vector autoregressions. J. Bus. Econom. Statist. 40 (3), 1268–1281. http://dx.doi.org/10.1080/07350015.2021.1920962.

Montiel Olea, J.L., Pflueger, C., 2013. A robust test for weak instruments. J. Bus. Econom. Statist. 31 (3), 358–369. http://dx.doi.org/10.1080/00401706.2013.806694.

Montiel Olea, J.L., Plagborg-Møller, M., Qian, E., Wolf, C.K., 2024. Double robustness of local projections and some unpleasant VARithmetic. Working Paper 32495, National Bureau of Economic Research, http://dx.doi.org/10.3386/w32495.

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

Nakamura, E., Steinsson, J., 2018. High-frequency identification of monetary non-neutrality: The information effect. Q. J. Econ. 133 (3), 1283–1330. http://dx.doi.org/10.1093/qje/qjy004.

Piger, J., Stockwell, T., 2025. Differences from differencing: Should local projections with observed shocks be estimated in levels or differences? J. Appl. Econometrics http://dx.doi.org/10.1002/jae.70003.

Rigobon, R., Sack, B., 2004. The impact of monetary policy on asset prices. J. Monet. Econ. 51 (8), 1553–1575. http://dx.doi.org/10.1016/j.jmoneco.2004.02.004.

Schlaak, T., Rieth, M., Podstawski, M., 2023. Monetary policy, external instruments, and heteroskedasticity. Quant. Econ. 14 (1), 161–200, doi:doi.org/10.3982/QE1511.

Sims, C.A., 1980. Macroeconomics and reality. Econometrica 48 (1), 1–48.

Stock, J.H., Watson, M.W., 2001. Vector autoregressions. J. Econ. Perspect. 15 (4), 101–115. http://dx.doi.org/10.1257/jep.15.4.101.

Swanson, E.T., 2021. Measuring the effects of Federal Reserve forward guidance and asset purchases on financial markets. J. Monet. Econ. 118, 32–53. http://dx.doi.org/10.1016/j.jmoneco.2020.09.003.

Uhlig, H., 2005. What are the effects of monetary policy on output? Results from an agnostic identification procedure. J. Monet. Econ. 52 (2), 381–419. http://dx.doi.org/10.1016/j.jmoneco.2004.05.007.

Wright, J.H., 2012. What does monetary policy do to long-term interest rates at the zero lower bound? Econ. J. 122 (564), F447–F466.