Shockingly Simple Monetary Shocks

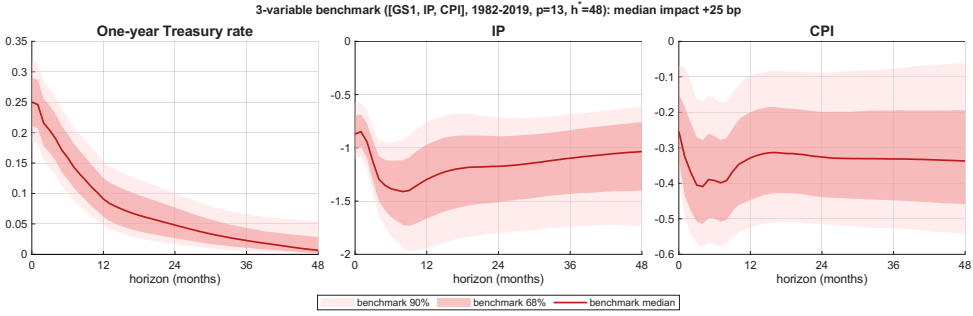
Tom D. Holden Mu-Chun Wang

Deutsche Bundesbank

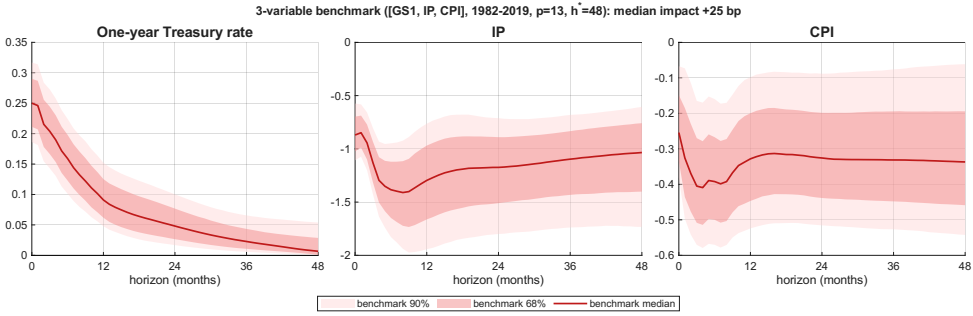
November 2026

The views expressed in this paper are those of the authors and do not represent the views of the Deutsche Bundesbank, the Eurosystem or its staff.

What kind of shock is this? Rate ↑, Output ↓, Prices ↓



What kind of shock is this? Rate ↑, Output ↓, Prices ↓



Produced with two simple restrictions:
Rates are positive and decline monotonically following monetary shocks.

Justifications for this restriction

1. A monetary shock is a mistake. If you make a mistake, you try to reverse it.
2. 3 eq. NK optimal policy satisfies the restriction after an inflation aversion shock.
3. Responses in estimated medium scale DSGE models ($\approx$) satisfy this restriction.
4. Responses in high-frequency identified monetary VARs ($\approx$) satisfy this restriction.

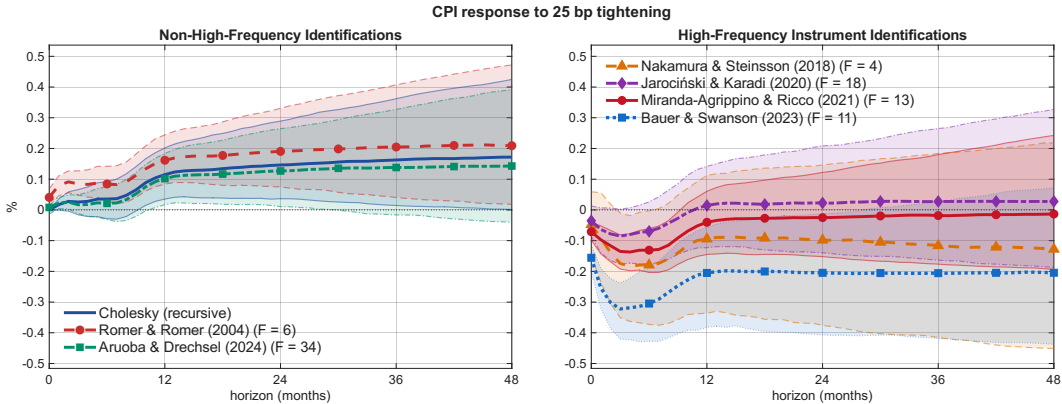
More details to come.

How the profession identifies monetary shocks

- **High-frequency**: Gertler Karadi (2015), Nakamura Steinsson (2018), Jarociński Karadi (2020), Miranda-Agrippino Ricco (2021), Bauer Swanson (2023).
- **Narrative**: Romer Romer (2004).
- **Text / NLP**: Aruoba Drechsel (2024).

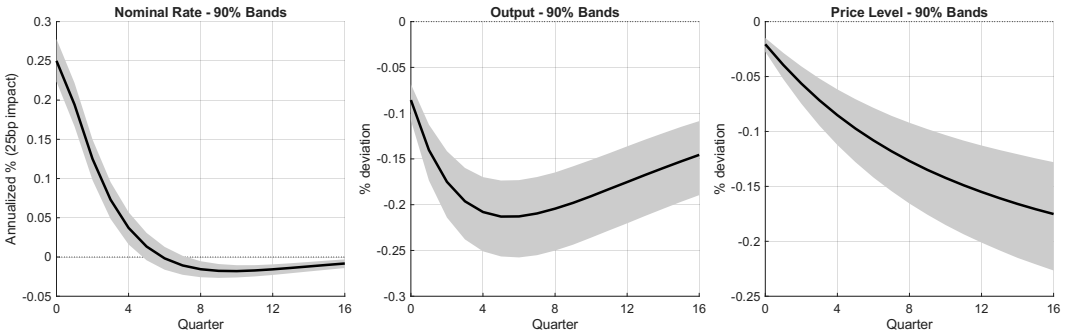
- Powerful: But only where a valid, strong instrument exists.

But the instruments disagree, and many are weak



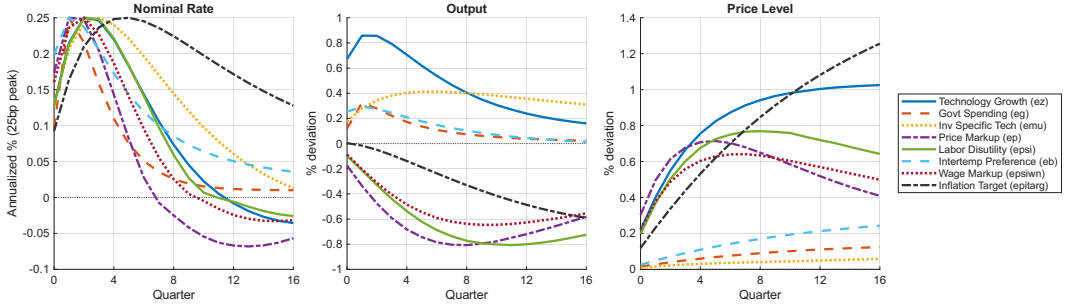
4 variable VAR including EBP. 1982-2019. Left: Cholesky, RR, AD. Clear price puzzles. Disagreement on sign and size ($+0.2\%$ to -0.3%). Many weak instruments (NS $F \approx 4$, RR $F \approx 6$).

How does the policy rate respond to monetary shocks in estimated medium-scale DSGE models?



IRFs to a monetary shock in a modified version of the Justiniano, Primiceri, and Tambalotti (JPT) (2013) model.

How does the policy rate respond to other shocks in estimated medium-scale DSGE models?



IRFs to non-monetary shocks in the modified JPT (2013) model.
Non-monetary shocks generate hump-shaped rate dynamics.

Model details: JPT (2013) + ...

- Estimated on 1964:Q3-2019:Q4 with GS1 instead of FFR,
- With hours in differences, and wages in double differences,
- + with the following model changes, all of which improve marginal data density:
  1. Canova (2014) style measurement errors (IID+I(1) on all, I(2) on wages.)
  2. Lomax shock priors (tails matched to original priors, second param KL matched).
  3. Central bank sets GS1, not FFR.
  4. Corrected entry of the inflation target in the monetary rule ($\mathbb{E}_t \pi^*_{t+1}$ enters).
  5. Estimated inflation target persistence.
  6. Identical priors on ϕ_π and ϕ_y.
  7. AR(2) monetary rule, with unit coefficient on first lag.
  8. Current quarter variables in the monetary rule.
  9. No price indexation, so price dispersion enters.
  10. No wage indexation, so wage dispersion enters.
  11. Automatic response of government spending to output growth.
  12. AR(2) government spending, with unit coefficient on first lag.
  13. Partially anticipated government spending (MA(1)).

What is the optimal response to a monetary policy shock?

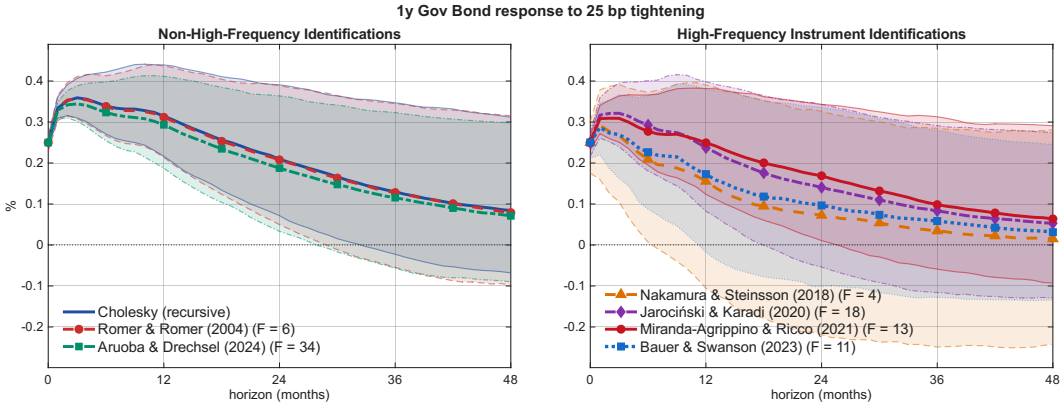
Model a monetary shock as one-period rise in the central bank's distaste for inflation. Objective: $\mu\pi_0 + \frac{1}{2}\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\pi_t^2 + \lambda x_t^2\right)$. Commitment. Rest of model: Basic NK.

- To cut today's inflation the bank optimally uses both tools: Lower expected inflation and open an output gap.
- It commits to a persistent disinflation and recession.

- Nominal rates = real rates ($\uparrow$) + expected inflation ($\downarrow$).
- Under plausible calibrations, the real rate rise dominates.
- $\Rightarrow$ Nominal rates rise. One (co-)state $\Rightarrow$ monotonic decline.

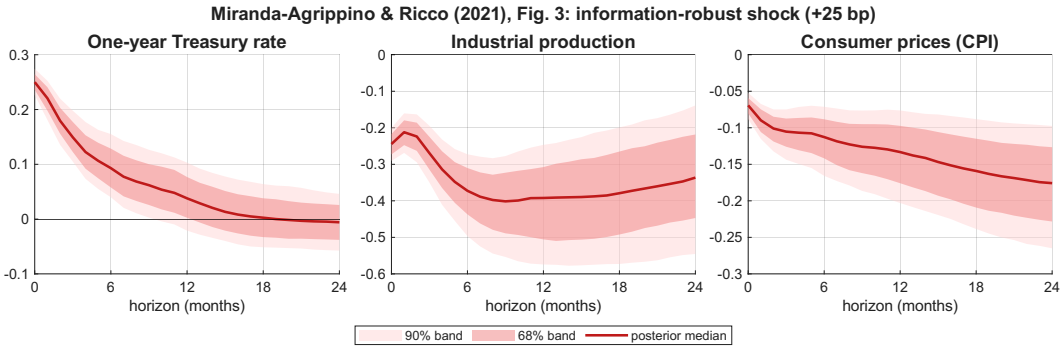
So, the optimal response to a monetary shock features a nominal interest rate that: **Rises on impact**, **declines monotonically** (no hump), **stays positive** (no overshoot).

Published monetary shock instruments violate this shape



Small humps ubiquitous. Some have substantial mass below zero.

Representative high-frequency responses



In Miranda-Agrippino & Ricco's (2021) own spec., their instrument gives clean results: Rate declines monotonically, output contracts, prices fall with no price puzzle.

Their Figure 3 spec., information-robust shock. 1979:1–2014:12. $F \approx 15$.

Our restriction: Monotone decay + non-negativity

Let q be the first column of the rotation matrix.
We can find vectors $a_0, a_1, \ldots$ such that the nominal rate IRF at horizon h is $a_h^\top q$.
Then, given a fixed cutoff h^*, we impose:

(i)	$a_0^\top q \geq 0$	rises on impact
(ii)	$-\left(a_h - a_{h-1}\right)^\top q \geq 0$	$h = 1, \ldots, h^*$
(iii)	$a_{h^*}^\top q \geq 0$	non-negative at h^*

Linear in q. $\Rightarrow$ The admissible rotations form a convex spherical polytope.

Our model, prior, and sampler

Villani (2009) style approach for dealing with the trend:

$$y_t = z_t + b\,(t - \bar{t}), \qquad z_t = c + \sum_{\ell=1}^{p} A_\ell z_{t-\ell} + u_t, \qquad b_{\text{rate}} = 0,$$

Joint prior:

$$\pi(b, B, \Sigma) \propto \mathcal{N}(b \mid \mu_{b0}, V_{b0}) \times \text{NIW}_\lambda(B, \Sigma \mid z(b)) \times \mathbf{1}[\{q \mid C_{B,\Sigma} q \geq 0, \|q\| = 1\} \neq \emptyset].$$

- NIW_λ uses $\alpha = 2$, MAR (2021) plug-in $\psi_i = \hat{\sigma}^2_{i,\text{AR}(1)}$ & GLP (2015) tightness $\hat{\lambda}$.
- Prior on b from pre-sample, with HAC covariance. Uninformative prior on c.
- We alternate pCN-MH for (B, Σ) with Gibbs for (b, c), and draw q uniformly.

Difference to Arias et al. (2018)

- Arias et al. (2018) implicitly imposes a volume-weighted prior on (B, Σ).
- Values of (B, Σ) for which the identified set has small measure get downweighted.

- Our conditionally-uniform prior minimises distortion of reduced form posterior.
- Partial answer to Baumeister and Hamilton (2015).

> *"we were able to conduct inference based on a reference prior that is conditionally uniformly distributed on the (compact) identified set $\Theta(\phi)$. Such a prior might have appeal to a broad audience."*
>
> — *Moon and Schorfheide (2012)*

- Slight caveat: We are uniform in "q-space", not in IRF space.

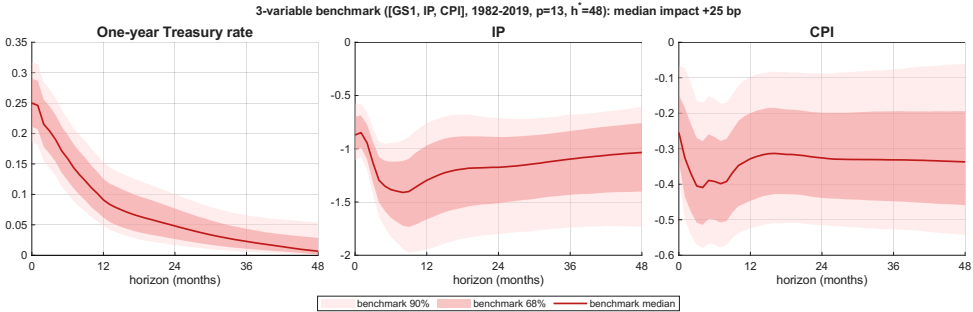
Data and sample

- Three monthly U.S. variables:
  1. one-year Treasury yield (%);
  2. industrial production ($100 \times$ log);
  3. consumer price index ($100 \times$ log).

- Sample **1982:01–2019:12** (post-Volcker), **13 lags** ($T_{\text{eff}} = 443$).
- Pre-sample (for b) **1947:01–1981:12**.

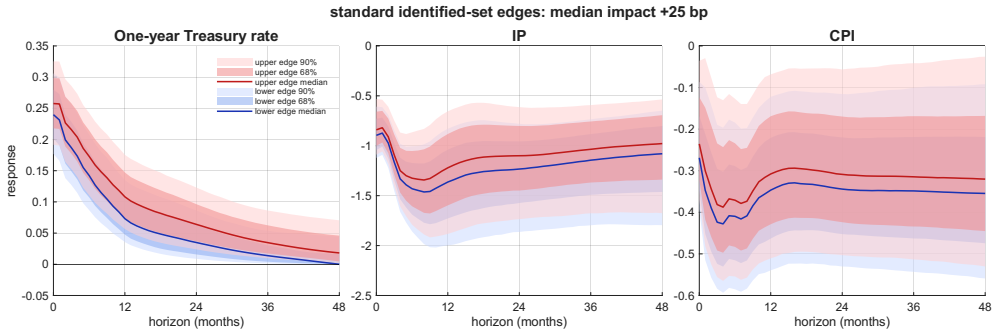
- Overall tightness $\hat{\lambda} = 0.38$ at the GLP marginal-likelihood mode.
- Restriction horizon $h^* = 48$ months, imposed on the one-year yield.
- Normalised to a $+25$ bp **median** impact rise in the rate.

Main result: A textbook tightening



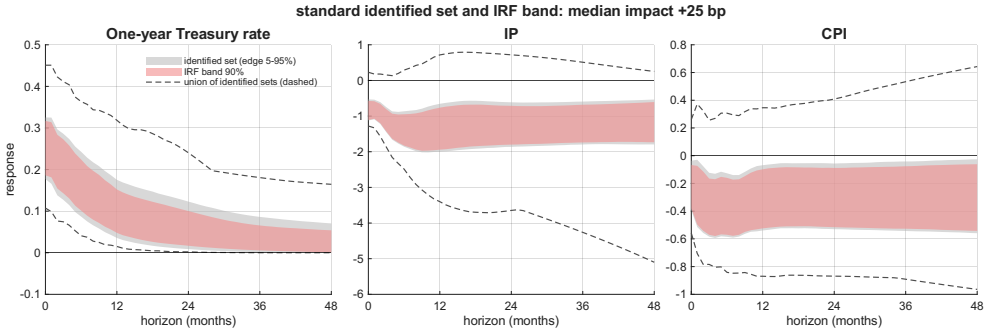
Rate rises and decays with no overshoot. Output contracts to a trough at $\sim$1 year.
The price level falls rapidly and permanently with no price puzzle.
No external instrument needed!

Edges of the identified set



Credible bands for the lower (blue) and upper (red) edges of the identified set.
Almost all the width reflects reduced-form uncertainty about where the set sits.
Near point identification! (Note: Edges are exactly point identified.)

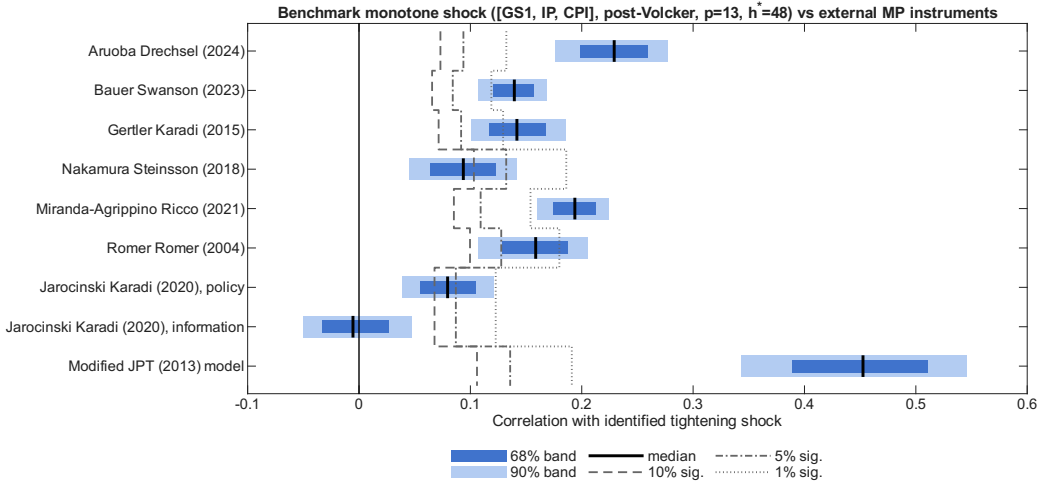
Union of identified sets



The dashed lines show the union of identified sets across all draws. Unsurprisingly wide.

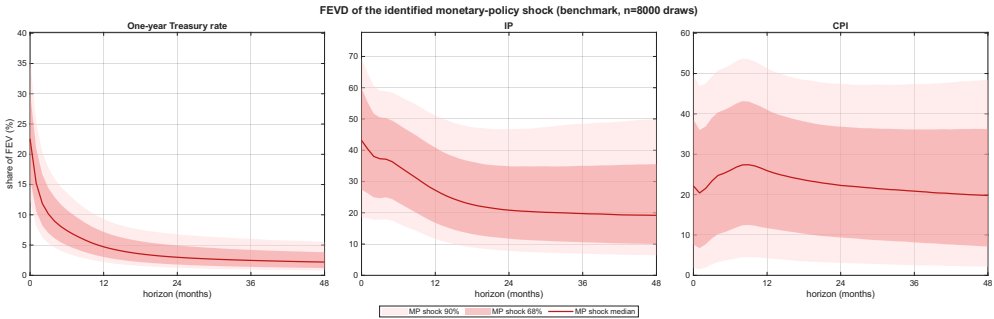
The grey area repeats the shaded area from the last slide.
The red area repeats the 90% credible bands from the main IRFs.
$\Rightarrow$ Baseline credible bands $\approx$ Credible bands of the upper & lower edges.

Validation versus other shocks



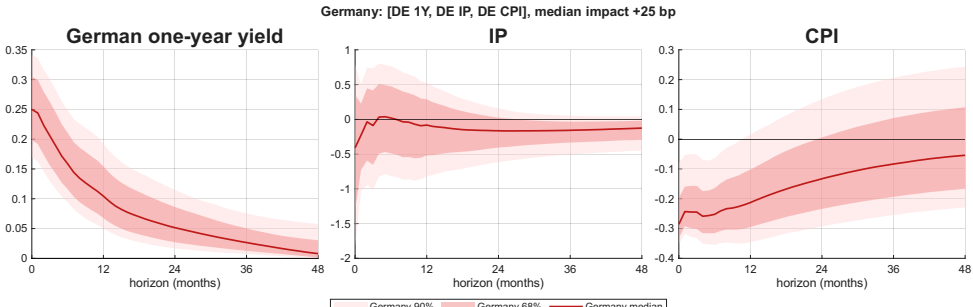
Correlates positively with all eight shocks. ≈ 0 corr with CB *information* shock, as desired.

Forecast error variance decomposition



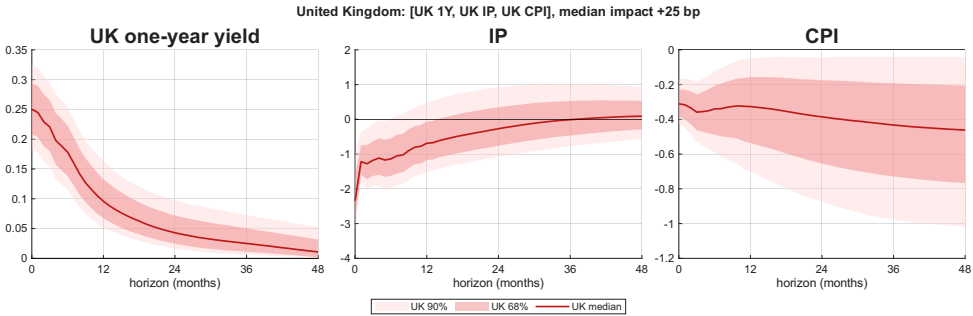
Share of forecast error variance explained by the identified monetary shock. Rather high, but this is inevitable in a small VAR. 68% and 90% credible bands.

Robustness: Germany



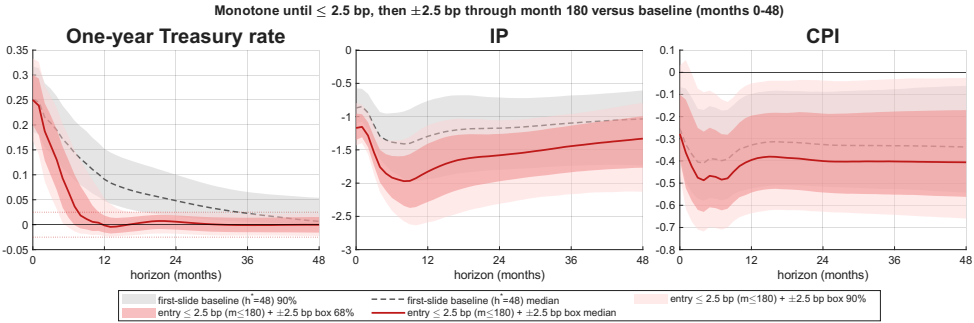
Germany [1y Bund yield, German IP, German CPI], 1982:01–2019:12 ($\hat{\lambda} = 0.28$).
Prices fall significantly in year 1; output contracts at medium horizons.

Robustness: United Kingdom



UK [1y gilt yield, UK IP, UK CPI], 1992:10–2019:12 ($\hat{\lambda} = 0.38$).
Post-ERM floating regime: output contracts in year 1; prices fall across all four years.

Alternative “box” identification: Decline then settle near zero

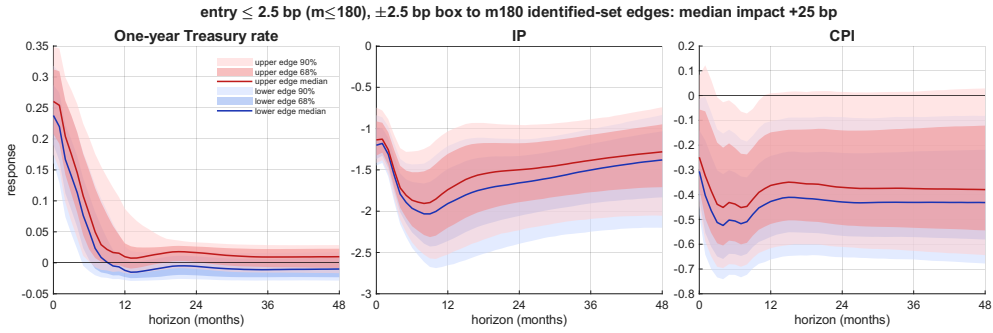


Benchmark (grey) vs new constraints (red):

Decreasing till $\leq 10\%$ of initial response, then stay within $\pm 10\%$ of initial response.

Faster rate response, and larger IP response.

Alternative “box” identification: Identified set edges



Credible bands (68% and 90%) for the upper and lower edges of the identified set.

Monotone identification in local projections

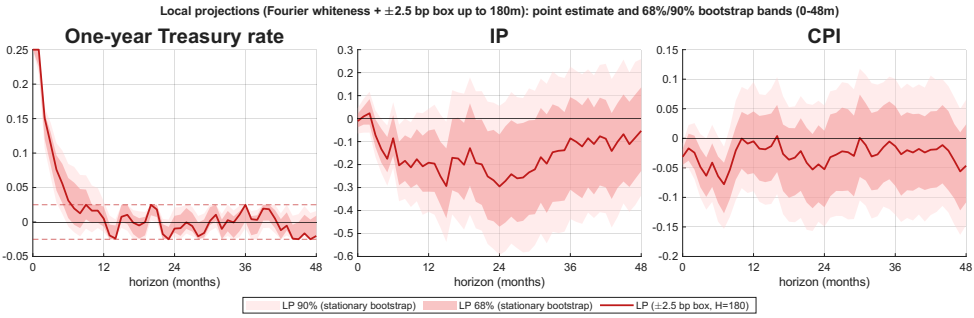
- Linear shape constraints leave the LP shock $e \in \mathbb{R}^{T-L}$ underdetermined.
- Find the maximally white shock by minimizing squared sample autocorrelations:

$$\sum_{j=-(T-L-1)}^{T-L-1} \gamma_j(e)^2,$$

- Subject to orthogonality of shocks to: a constant; a trend; 13 lags of GS1, IP, CPI.
- And subject to settle-near-zero (± 2.5 bp) box restriction:

$$0 \le a_h \le a_{h-1} \quad (1 \le h < k^*), \qquad |a_h| \le 2.5 \text{ bp} \quad (k^* \le h \le 180).$$

Local projections (± 2.5 bp box)



LP point estimate and 68%/90% stationary bootstrap bands:
Both IP and CPI contract with no price puzzle.
Posterior median corr. of LP shocks with benchmark ones is 0.48 (90%: $[0.37, 0.57]$).

Conclusion

- Simple shape restrictions on policy-rate IRF: non-negative, monotonic decline.
- Delivers a textbook monetary tightening in a minimal three-variable VAR.

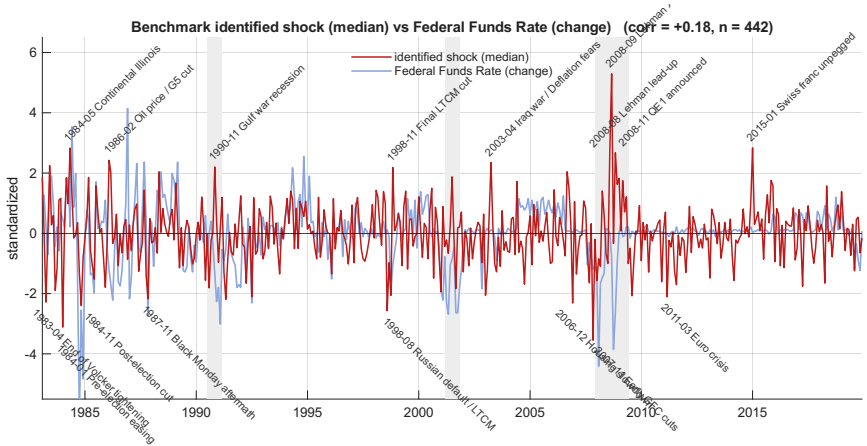
- No high-freq. info, no sign restriction on other variables, no recursive ordering.
- Correlates with all eight monetary-policy shocks, *not* the information shock.

- Robust to horizon, sample, prices, prior, larger VARs, countries, box, and LPs.

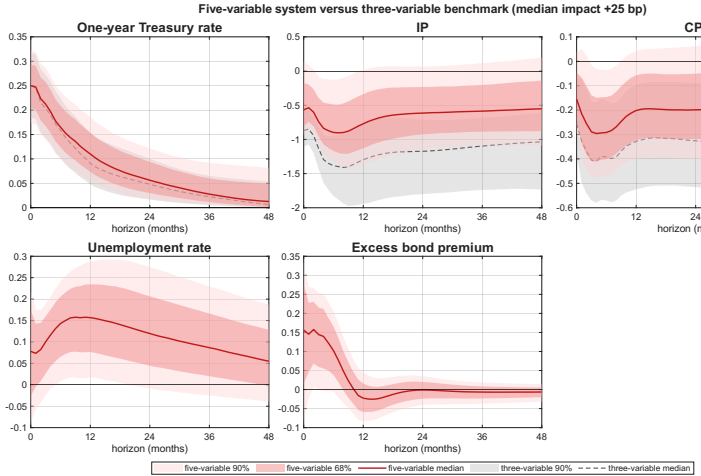
- Rapid response of prices to monetary shocks is not an artefact of HF identification.

Comments very welcome!

The shocks

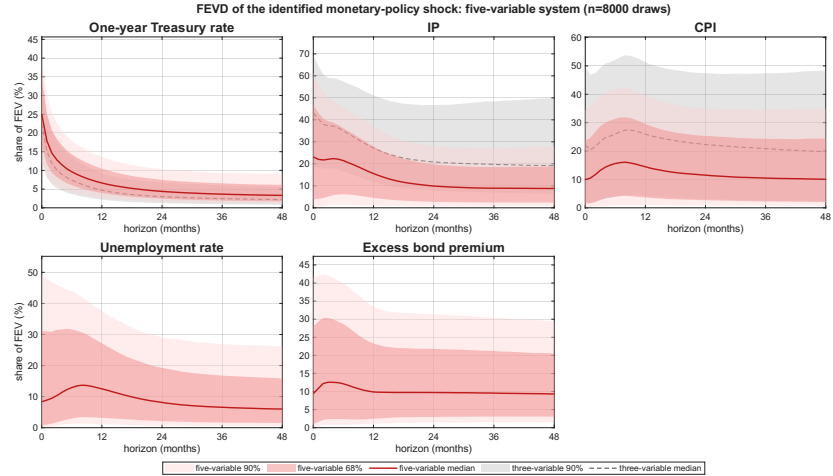

Benchmark identified shock (median) vs Federal Funds Rate (change) (corr = +0.18, n = 442)
identified shock (median)
Federal Funds Rate (change)
standardized
1984-05 Continental Illinois
1986-02 Oil price / G5 cut
1990-11 Gulf war recession
1998-11 Final LTCM cut
2003-04 Iraq war / Deflation fears
2008-09 Lehman /
2008-08 Lehman lead-up
2008-11 QE1 announced
2015-01 Swiss franc unpegged
1983-04 Fed / Volcker tightening
1984-01 Pre-election easing
1984-11 Post-election cut
1987-11 Black Monday aftermath
1998-08 Russian default / LTCM
2006-12 Housing / Fed tightening
2007-12 Financial crisis / ECB cuts
2011-03 Euro crisis

Five-variable VAR: Adding unemployment and EBP



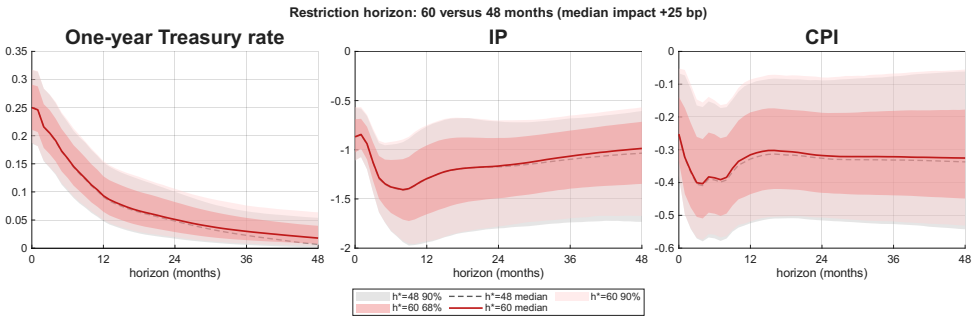
Adding unemployment and the excess bond premium (red, $\hat{\lambda} = 0.25$) vs the 3-var benchmark (grey).

Forecast error variance decomposition: Five-variable VAR



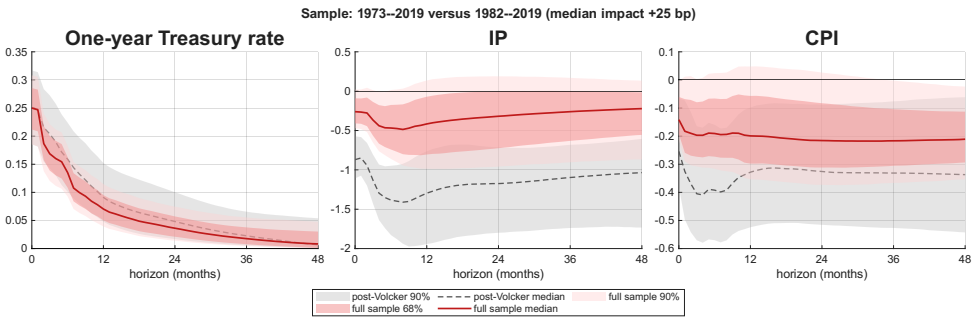
Explains 10–15% of 1-year IP, CPI, unemployment, and EBP variance (red), vs 6–10% at medium horizons.

Robustness: Restriction horizon



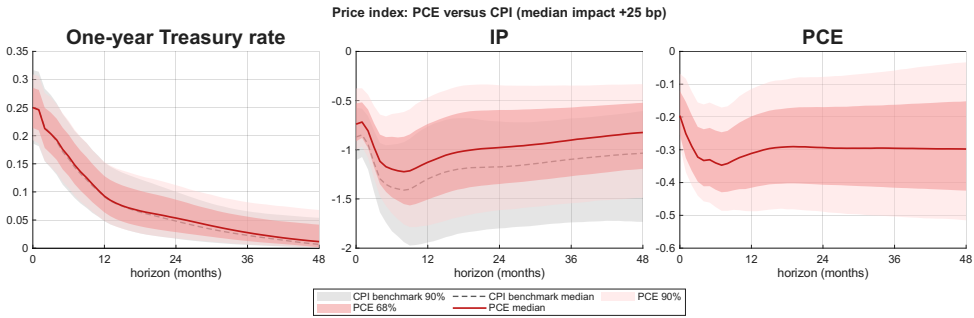
Benchmark $h^* = 48$ (grey) vs $h^* = 60$ (red). The picture is unchanged.

Robustness: Sample period



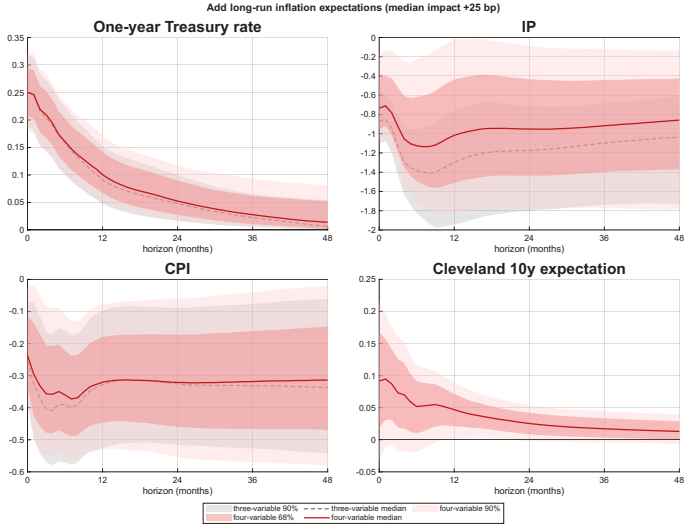
Post-Volcker benchmark 1982–2019 (grey) vs full sample 1973–2019 (red): Similar rate dynamics, slightly weaker responses.

Robustness: Price index



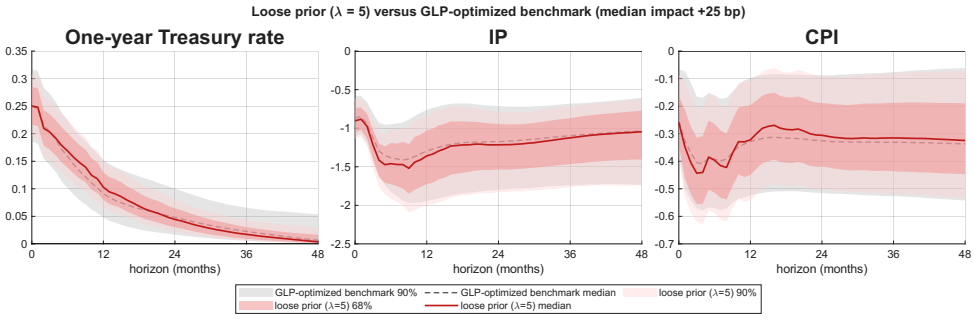
CPI benchmark (grey) vs the PCE price index (red). No price puzzle under either.

Robustness: Adding long-run inflation expectations



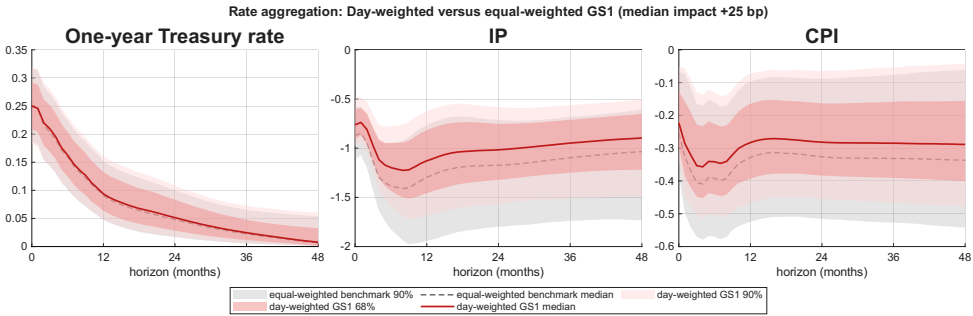
Adding the Cleveland 10y expected-inflation series (red) vs the 3-var benchmark (grey).

Robustness: Prior tightness



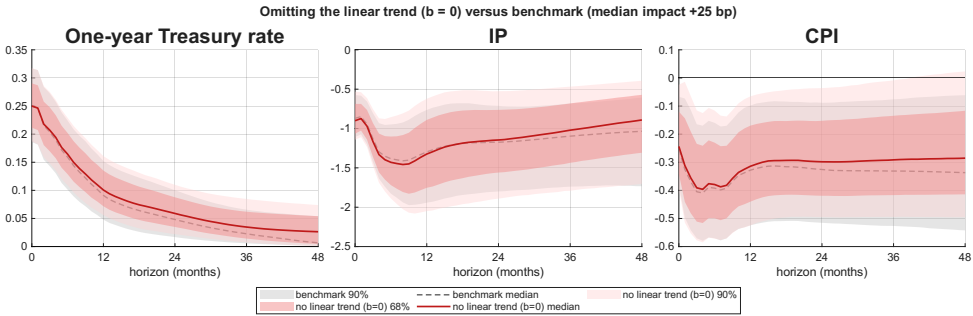
Loose Minnesota prior $\lambda = 5$ (red) vs the GLP-optimized benchmark $\hat{\lambda} = 0.38$ (grey).

Robustness: Day-of-month weighted interest rate



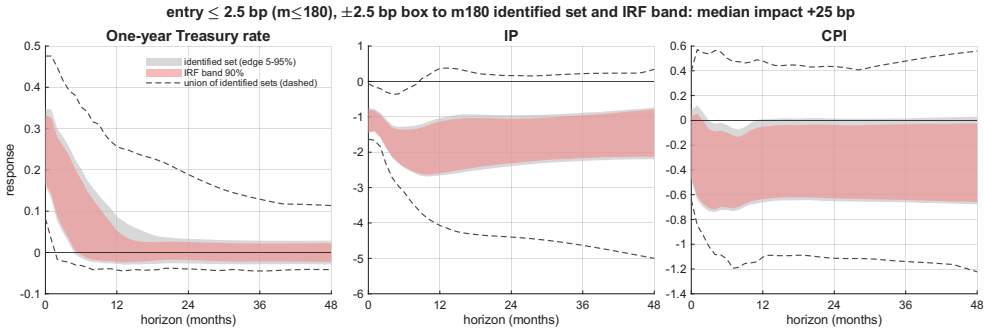
Weighting daily `GS1` by calendar day $d \in \{1, \ldots, 31\}$ (red) vs flat average (grey): Tilts toward post-FOMC information without last-day noise.

Robustness: Omitting the linear trend ($b = 0$)



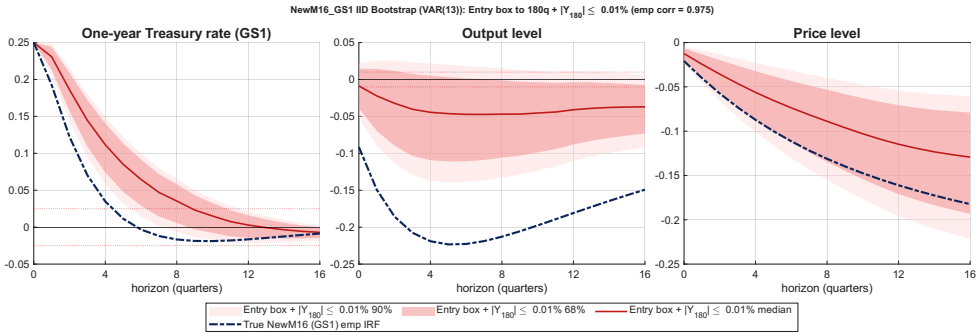
Omitting the linear trend ($b = 0$, red) vs the linear-trend benchmark (grey): Trend sharpens the medium-horizon price decline.

Alternative “box” identification: Identified set envelope



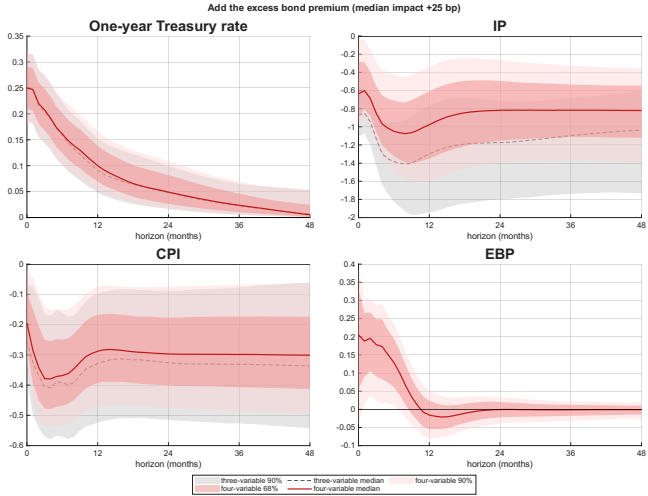
Headline 90% band (red) vs. identified set (grey) vs. union envelope (dashed).

Alternative "box" identification: Simulated DSGE data



VAR(13) on data simulated from modified JPT (2013) model.
80k quarters; ± 2.5 bp rate box to $q = 180$, $|Y_{180}| \leq 0.01\%$.
Posterior median corr. with true shock: 0.975 **with** $|Y_{180}| \leq 0.01\%$ vs 0.899 **without**.

Appendix: Adding the excess bond premium



Adding the Gilchrist–Zakrajšek (2012) excess bond premium (red, $\hat{\lambda} = 0.30$) vs the 3-var benchmark (grey).

Appendix: What +25 bp means in these results

Let $R_i^{(d)}(h)$ be the unnormalised response in posterior draw d. We use one positive factor for all draws and all variables:

$$s = \frac{0.25}{\text{median}_d\, R_{\text{GS1}}^{(d)}(0)}, \qquad \widetilde{R}_i^{(d)}(h) = sR_i^{(d)}(h).$$

- The median rate impact is +25 bp; individual impacts still vary.
- We calculate medians and credible bands *after* applying this common factor.
- Since $s > 0$, normalisation cannot reverse a draw's sign.
- Each alternative specification receives its own median-impact scale.

Appendix: The pCN-MH / Elliptical Slice Sampler

Block updates target $\pi(b, B, \Sigma) \propto \mathcal{N}(b \mid \mu_{b0}, V_{b0})\,\mathrm{NIW}_\lambda(B, \Sigma \mid z(b))\,\mathbf{1}[\mathcal{S}(B, \Sigma) \neq \varnothing]$ $\big(L = \mathrm{chol}\,\Sigma,\ R = \mathrm{chol}\,Q(b)^{-1}\big)$:

B — **preconditioned Crank–Nicolson / Elliptical Slice Sampling (ESS):**

$$B(\theta) = \bar{B}(b) + \cos\theta\,\big(B - \bar{B}(b)\big) + \sin\theta\,R\,Z\,L', \qquad Z_{ij} \overset{\text{iid}}{\sim} \mathcal{N}(0,1).$$

Great-circle rotations preserve $\mathcal{MN}\big(\bar{B}(b), \Sigma \otimes Q(b)^{-1}\big)$, so acceptance reduces to checking $\mathcal{S}(B(\theta), \Sigma) \neq \varnothing$ (θ chosen by windowed bracket shrinkage; subspace steps on $\sum_\ell A_\ell$).

Σ — **conditional Inverse-Wishart Gibbs + affine-invariant MH** (every 5th step): exact draw $\Sigma' \sim \mathcal{IW}\big(S^*(b) + \Delta B^\top Q(b) \Delta B,\ d^* + m\big)$ accepted iff $\mathcal{S}(B, \Sigma') \neq \varnothing$, paired with a Cholesky-symmetric MH step $L' = L \exp\big(\frac{\eta}{2}(W + W')\big)$.

b — **Gaussian Transport MH + block-Gaussian Gibbs** (every 5th step): transports $B' = \bar{B}(b') + R(b')\big(R(b)^{-1}(B - \bar{B}(b))\big)$ along the conditional ridge as b' moves.

Initialised at a feasible (B_0, Σ_0) found by rejection at a tight $\lambda_{\text{init}} = 0.1$, then run at the optimal $\hat{\lambda}$.

Appendix: Drawing the rotation: Geodesic hit-and-run

Each feasible (B, Σ) the pCN-MH chain retains still needs a rotation q (first column of Q): the shock's impact vector is Lq and the rate response at horizon h is $a_h^\top q$. The monotone-decay restriction is the spherical polytope $\mathcal{S} = \{q \in \mathbb{S}^{n-1} : Cq \geq 0\}$.

We draw q uniformly from $\mathcal{S}$ by geodesic hit-and-run:

1. at the current q, pick a random tangent direction d;
2. move along the great circle $\gamma(s) = \cos(s)\, q + \sin(s)\, d$;
3. each row $c_k^\top \gamma(s) \geq 0$ defines a feasible arc; intersect the K arcs;
4. draw s uniformly on the feasible interval, set $q \leftarrow \gamma(s)$, and repeat.

The polytope is geodesically convex, so the feasible set on each great circle is a single arc.
Chain converges to the uniform law on $\mathcal{S}$. We take one q per retained (B, Σ) (burn-in 200, thin 10).

Appendix: Stationary bootstrap for local projections

For variables $w \in \{i, y, p\}$, horizons $h \in \{0, \ldots, 180\}$, let $u_{w,h}$ be the OLS LP residual. (I.e. $u_{w,h}$ is the fitted residual from regressing w_h on the estimated shocks e, a constant, a trend and 13 lags of i, y, p over all but the last h periods.)

And let $\hat{\theta}_{w,h}$ be the estimated response of w at h with $\hat{\theta}_{i,0} = 1$.

In each of $B = 500$ Politis–Romano (1994) stationary bootstrap replications:

1. draw a circular block-resampled index path $\tau^{(b)}$ (geometric block length mean 12);
2. let $u^{(b)}_{w,h}$ be the residual from regressing $\tau^{(b)}$-shuffled $u_{w,h}$ on a constant, a trend, and 13 lags of $\tau^{(b)}$-shuffled i, y, p;
3. for all e, let $\hat{\theta}^{(b)}_{w,h}(e) = \hat{\theta}_{w,h} + \frac{1}{T-L-h}(u^{(b)}_{w,h})^\top e_{1:T-L-h}$.
4. solve for optimal $e^{(b)}$ subject to shape restrictions on $\hat{\theta}^{(b)}_{i,h}(e^{(b)})$; record $\hat{\theta}^{(b)}_{w,h}(e^{(b)})$.

Reported 68% and 90% bands are $[16, 84]$ and $[5, 95]$ percentiles across the $B = 500$ bootstrap draws.