Shockingly Simple Monetary Shocks

Tom D. Holden   Mu-Chun Wang

Deutsche Bundesbank

2026-09-07, Boston Fed

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

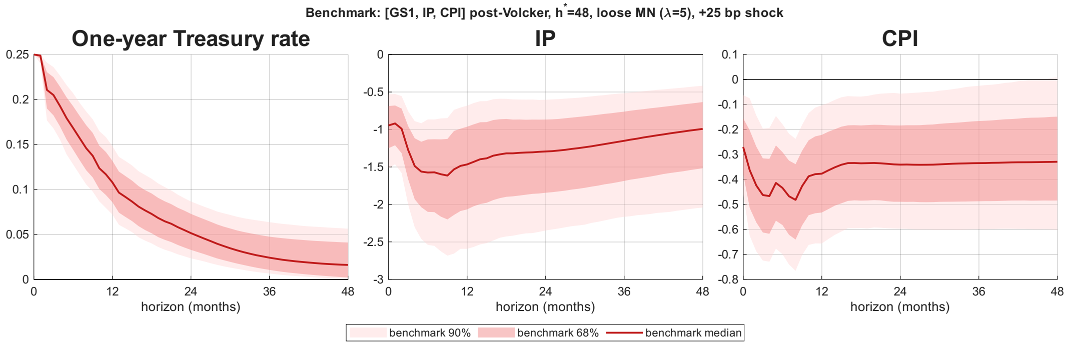

What kind of shock is this? Rate ↑, Output ↓, Prices ↓

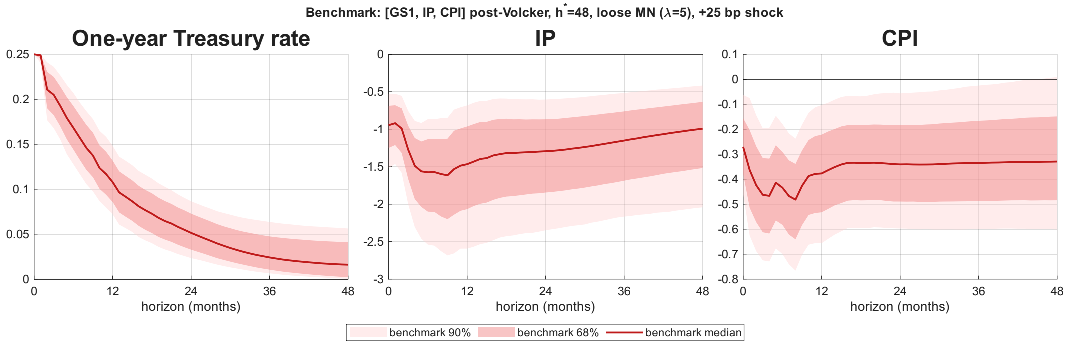

Produced with two simple restrictions:
Rates are positive and decline monotonically following monetary shocks.

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

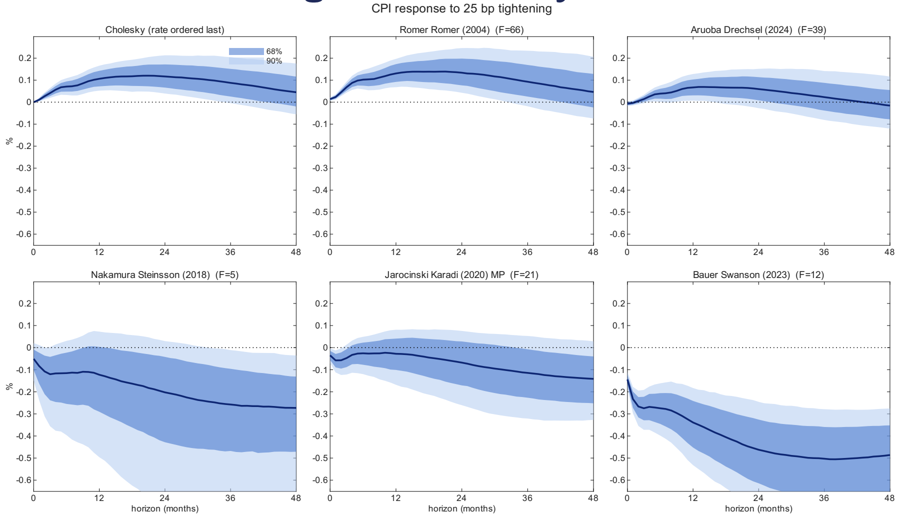

4 variable VAR including EBP. 1973-2019. Top row: Cholesky, RR, AD. Clear price puzzles. Disagreement on sign and size (+0.14% to −0.51%). Many weak instruments (NS $F \approx 5.2$).

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

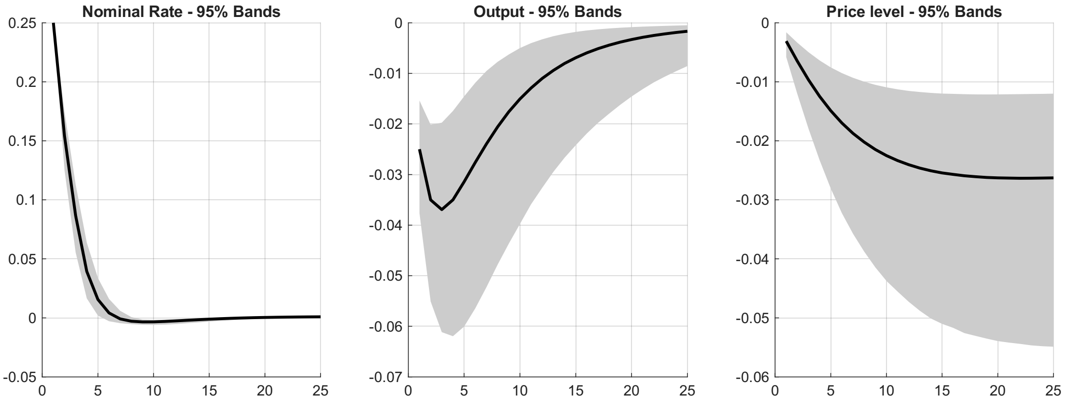

IRFs to a monetary shock in Justiniano, Primiceri, and Tambalotti (JPT) (2013).
(Minuscule overshoot (UB: -0.002).)

Improvements to the JPT model remove the overshoot

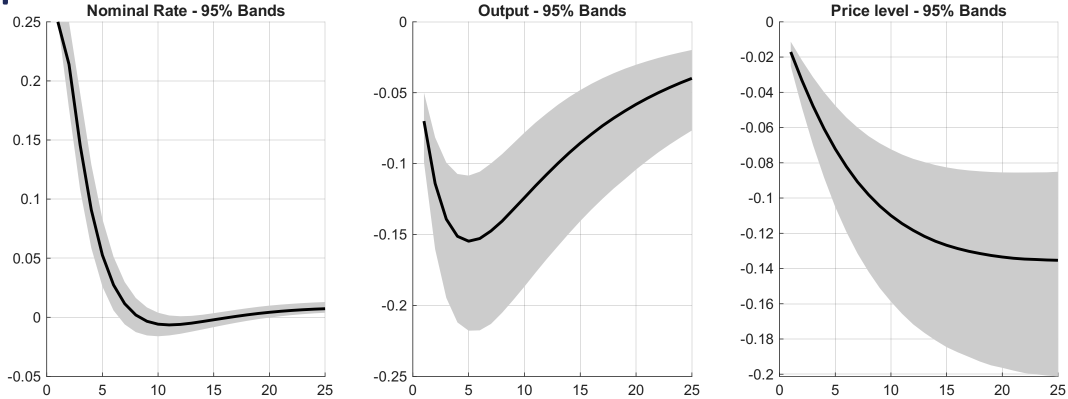

JPT Modified Harmonic Mean (MHM) Log Marginal Data Density (LMDD): -1138.54. Improvements:

- Remove certainly counter-factual price indexation: MHM LMDD: -1128.88
- Remove mostly counter-factual wage indexation: MHM LMDD: -1114.44
- Use current inflation in the monetary rule: MHM LMDD: -1103.62
- Estimate the monetary shock persistence: MHM LMDD: -1101.65

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: $\mu\pi_0 + \frac{1}{2}\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\pi_t^2 + \lambda x_t^2\right)$. 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

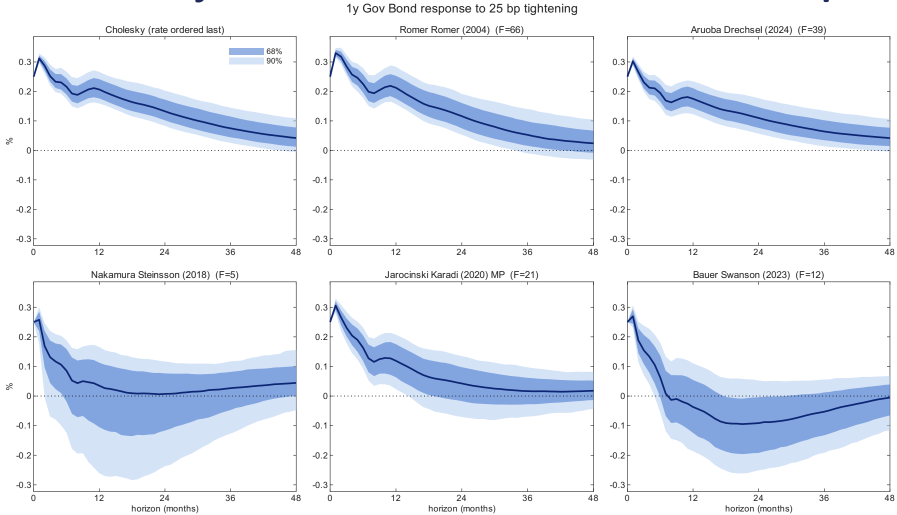

Same VAR, different instruments, very different *rate* responses: humps and substantial overshooting.
Across seven leading studies, 3/7 are hump-shaped, 4/7 overshoot below zero (in authors' IRFs).

What a well-identified shock looks like

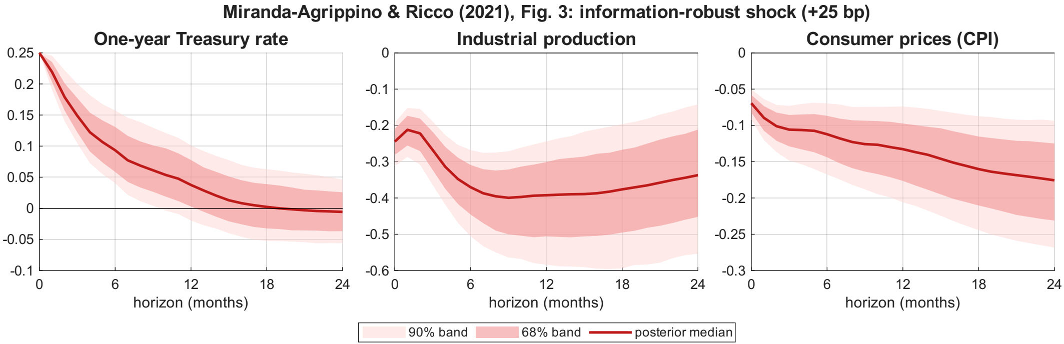

Miranda-Agrippino & Ricco's (2021) information-robust instrument gives clean results: Rate declines monotonically, output contracts, prices fall with no price puzzle.

(Their Figure 3 spec; information-robust shock; 68% and 90% posterior bands; +25 bp impact. 1979:1–2014:12. First-stage F-stat $F = 15.3$, reliability $\Lambda = 0.26$.)

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)	$-(a_h - a_{h-1})^\top q \geq 0$	$h = 1, \ldots, h^*$
(iii)	$a_{h^*}^\top q \geq 0$	non-negative at h^*

- Agnostic about decay *speed*, *magnitude*, *every other variable*.
- Linear in q.
- $\Rightarrow$ The admissible rotations form a convex spherical polytope.

Our prior and sampler

We use the following prior:

$$\pi(B, \Sigma) \propto \mathrm{NIW}_\lambda(B, \Sigma) \times \mathbf{1}[\{q \mid C_{B,\Sigma} q \geq 0, \|q\| = 1\} \neq \emptyset].$$

where $C_{B,\Sigma}$ collects the restrictions from the previous slide,
and NIW_λ uses the Giannone, Lenza, Primiceri (2015) version of the Minnesota prior.

We sample from (B, Σ) using a preconditioned Crank–Nicolson proposal.

Having sampled (B, Σ) from the posterior, we draw q uniformly from:

$$\{q \mid C_{B,\Sigma} q \geq 0, \|q\| = 1\}.$$

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), **12 lags** ($T = 444$).

- **Loose** Minnesota prior, $\lambda = 5$ (far above the usual ≈ 0.2).
- Restriction horizon $h^* = 48$ months, imposed on the one-year yield.
- Normalised to a $+25$ **bp** impact rise in the one-year yield.

Main result: A textbook tightening

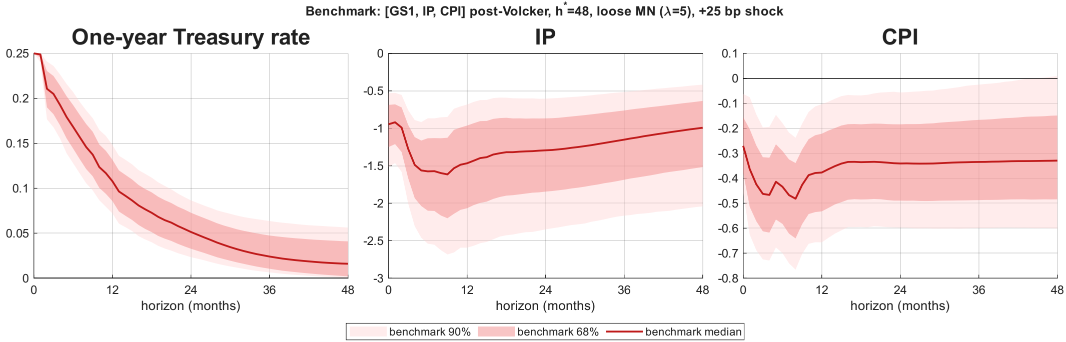

Rate rises and decays with no overshoot. Output contracts to a trough at ∼1 year.
The price level falls permanently with no price puzzle.
No external instrument needed!

Edges of the identified set

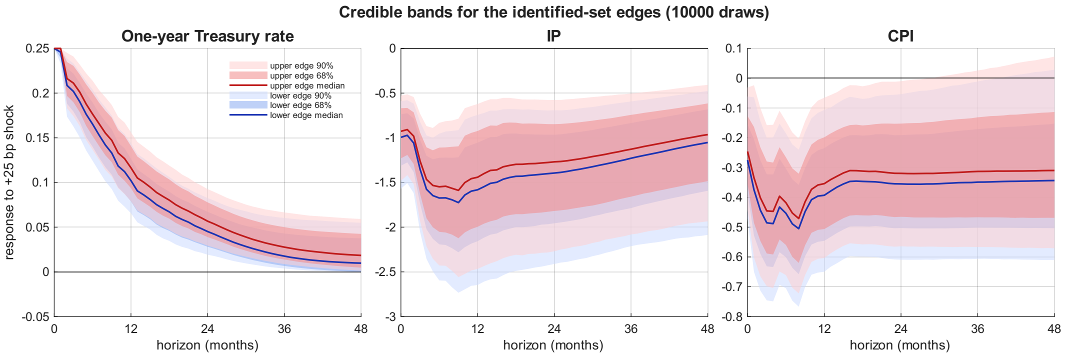

Credible bands for the lower (blue) and upper (red) edges of the identified set. Almost all the width reflects reduced-form uncertainty about where the set sits. Near point identification! (Note: Edges are exactly point identified.)

Union of identified sets

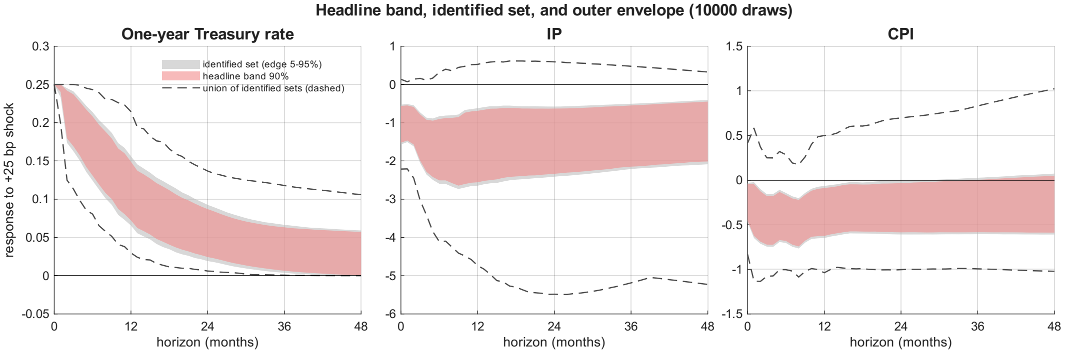

The dashed lines show the union of identified sets across all draws. Unsurprisingly wide.

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 instruments

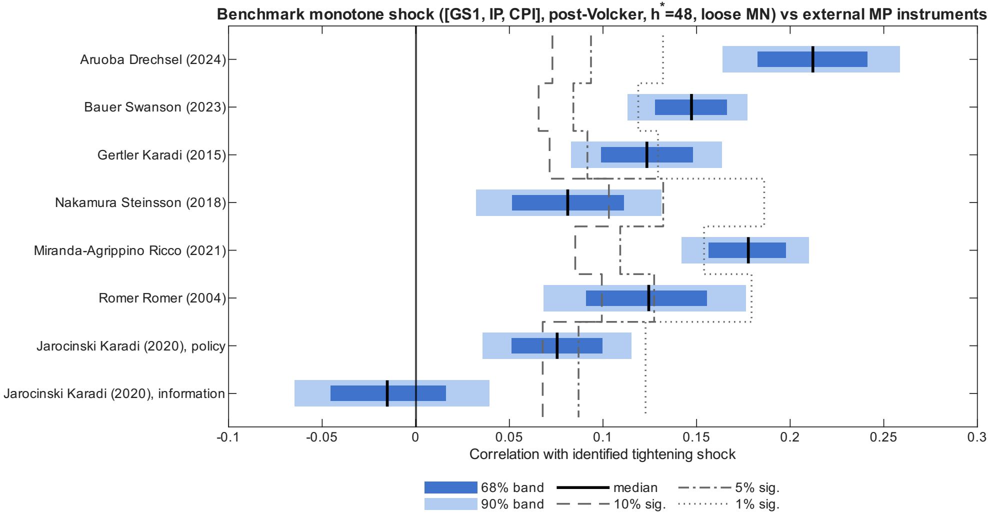

The shock uses *no* external information, yet it correlates positively with all seven monetary-policy instruments (90% band excludes zero) and is ≈ 0 with the CB *information* shock, as it should be.

The shocks

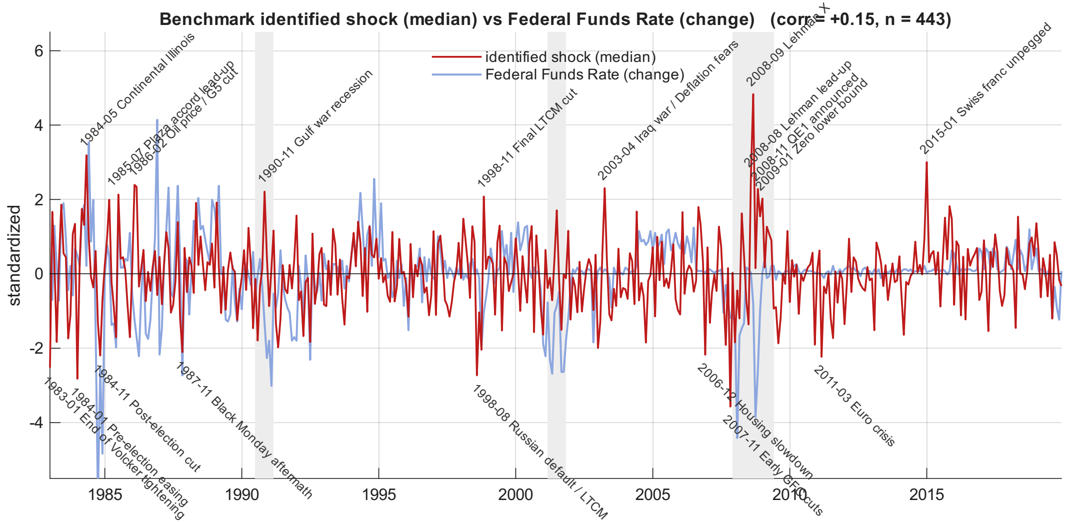
Benchmark identified shock (median) vs Federal Funds Rate (change) (corr = +0.15, n = 443)
identified shock (median)
Federal Funds Rate (change)
standardized
1984-05 Continental Illinois
1985-07 Plaza accord lead-up
1986-02 Oil price / G5 cut
1990-11 Gulf war recession
1998-11 Final LTCM cut
2003-04 Iraq war / Deflation fears
2008-09 Lehman X
2008-08 Lehman lead-up
2008-11 QE1 announced
2009-01 Zero lower bound
2015-01 Swiss franc unpegged
1983-01 End of 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 slowdown
2007-11 Early GFC cuts
2011-03 Euro crisis
1985
1990
1995
2000
2005
2010
2015
6
4
2
0
-2
-4

Robustness: Restriction horizon

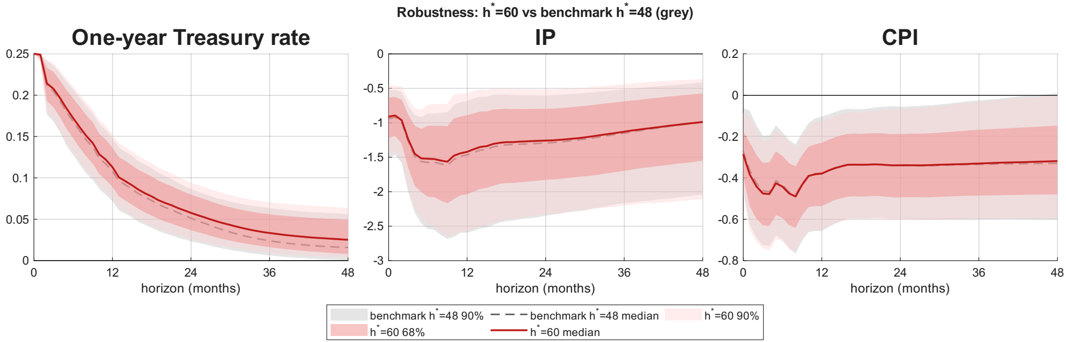

Benchmark $h^* = 48$ (grey) vs $h^* = 60$ (red) — the picture is unchanged.

Robustness: Sample period

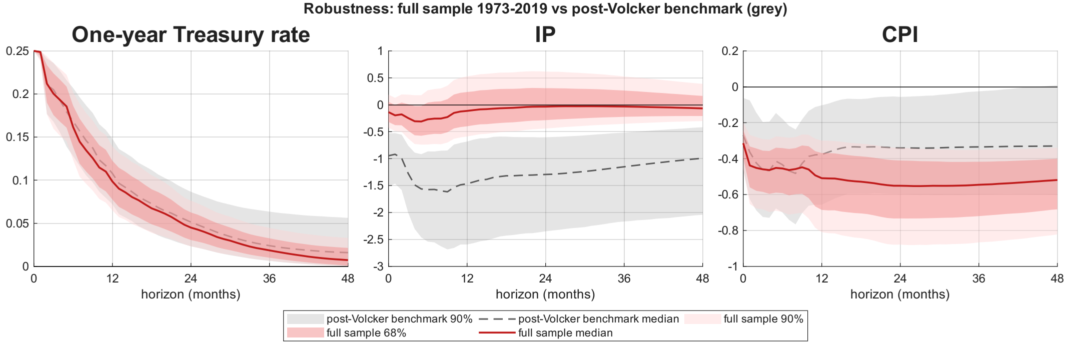

Post-Volcker benchmark 1982–2019 (grey) vs full sample 1973–2019 (red):
Similar rate and price dynamics. Weaker output response.

Robustness: Price index

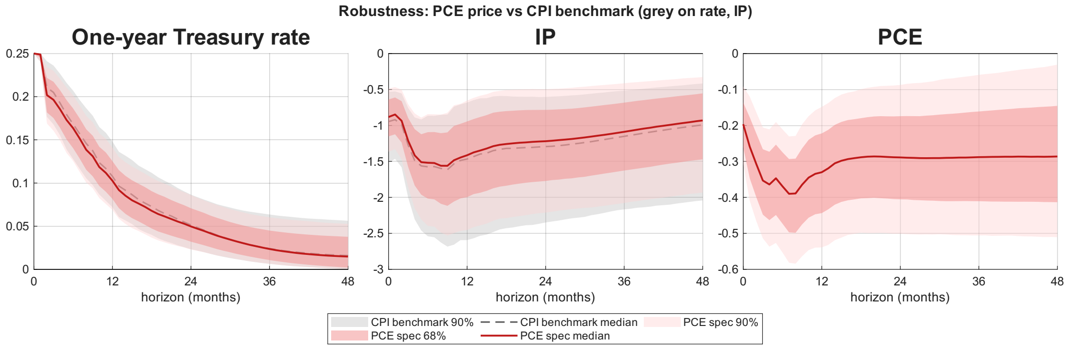

CPI benchmark (grey) vs the PCE price index (red) — no price puzzle under either.

Robustness: Adding an expectations channel

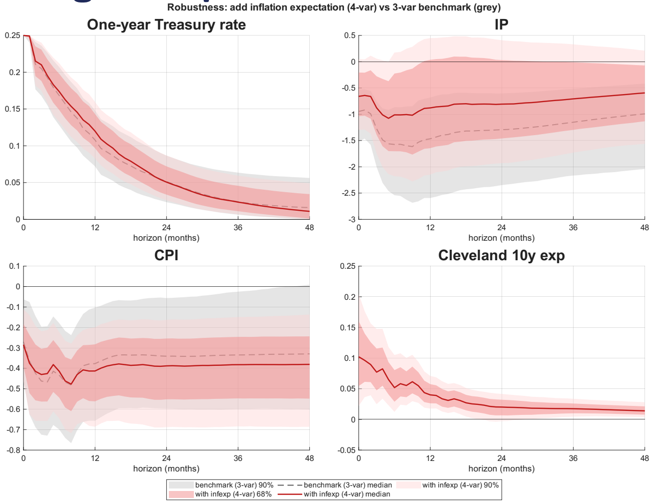

Adding the Cleveland 10y expected-inflation series (4-var, red) vs the 3-var benchmark (grey):
Rate, output, and price responses are essentially unchanged; long-horizon expectations stay anchored.

Robustness: Prior tightness

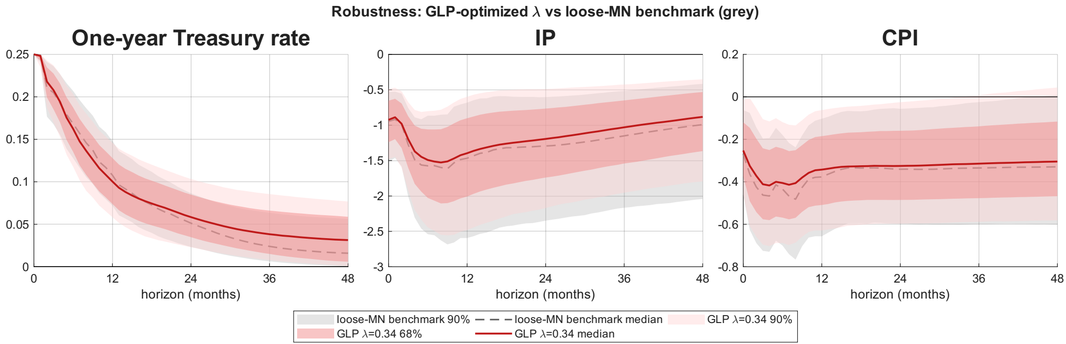

Loose benchmark $\lambda = 5$ (grey) versus...

the Giannone, Lenza, Primiceri (2015) hierarchical prior $\bar{\lambda} = 0.34$ (red).

Almost identical.

Robustness: The euro area

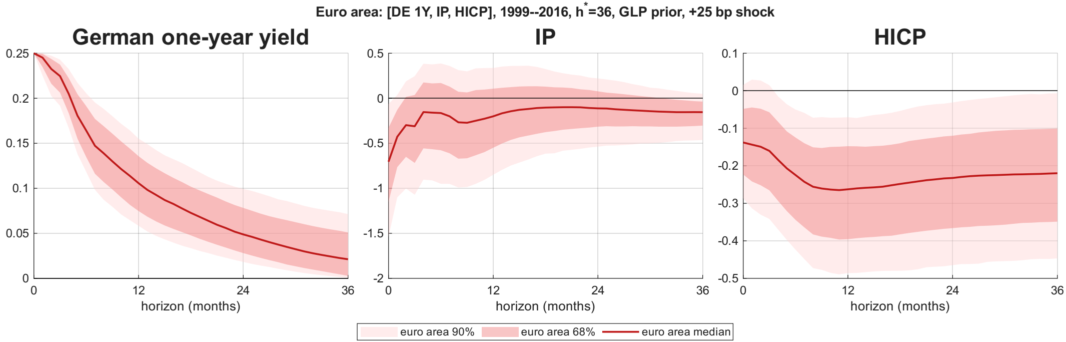

Euro area [German 1y yield, IP, HICP], 1999–2016, $h^* = 36$, same shape restriction.
Classic monetary tightening: HICP falls (no price puzzle) and output contracts.

Alternative "box" identification: Decline then settle near zero

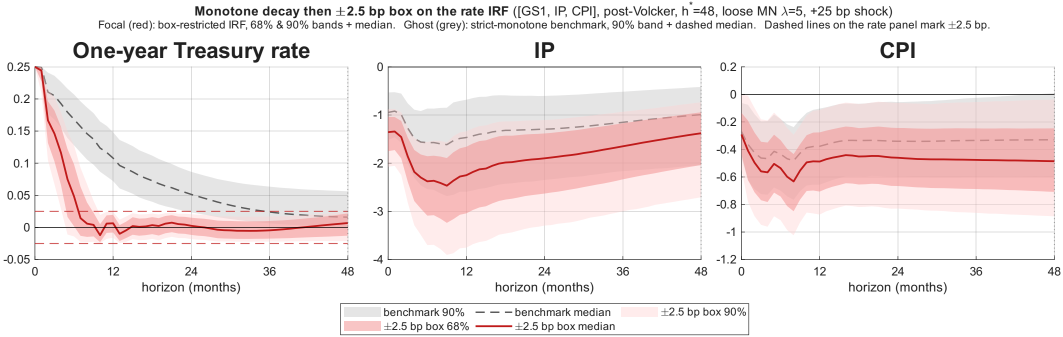

New constraints: Decay to zero, then stay within ±2.5 bp (red).
The variable responses closely track the benchmark ones (grey).

Alternative "box" identification: Identified set edges

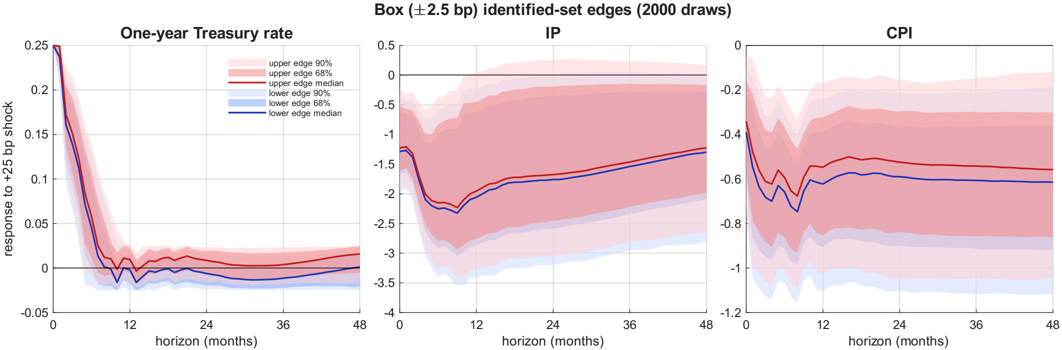

Credible bands (68% and 90%) for the upper and lower edges of the identified set.

Alternative "box" identification: Identified set envelope

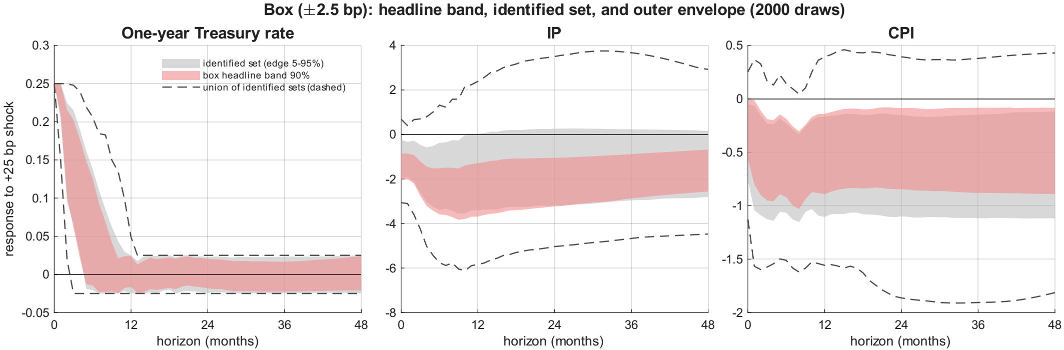

Headline 90% band (red) vs. identified set (grey) vs. union envelope (dashed).

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 seven monetary-policy instruments, *not* information shock.

- Robust to horizon, sample, price index, prior, expectations, settle-near-zero box.

Comments very welcome!

Appendix: The pCN-MH sampler

Two block updates target $\pi(B,\Sigma) \propto \mathrm{NIW}_\lambda(B,\Sigma)\,\mathbf{1}[\mathcal{S}(B,\Sigma) \neq \varnothing]$ $\big(L = \operatorname{chol}\Sigma$, $R = \operatorname{chol}(X'X)^{-1}\big)$:

B — **preconditioned Crank–Nicolson proposal:**

$$B' = \bar{B} + \sqrt{1-\beta^2}\,(B - \bar{B}) + \beta\, R\, Z\, L', \qquad Z_{ij} \overset{\text{iid}}{\sim} \mathcal{N}(0,1).$$

It is reversible w.r.t. the matrix-normal prior $\mathcal{MN}\big(\bar{B}, \Sigma \otimes (X'X)^{-1}\big)$, so the Metropolis ratio collapses to accept iff $\mathcal{S}(B',\Sigma) \neq \varnothing$ — a pure feasibility check, with *no* posterior density to evaluate. ($\beta = 0.10$, $\approx$45% accepted.)

Σ — **Cholesky-symmetric random walk** (every 5th step):

$$L' = L\,\exp\!\big(\tfrac{\eta}{2}(W + W')\big), \qquad \Sigma' = L' L'^{\top}.$$

The proposal is symmetric in the affine-invariant metric; the acceptance probability combines the NIW Metropolis ratio and a Jacobian $(n+1)\eta\,\mathrm{tr}(W)$. ($\eta = 0.02$, $\approx$30%.)

Initialised at a feasible (B_0, Σ_0) found by rejection at a tight $\lambda_{\text{init}} = 0.1$, then run at the loose $\lambda = 5$.

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 — the chain is well defined and converges to the uniform law on $\mathcal{S}$. We take one q per retained (B, Σ) (burn-in 200, thin 10); only inequalities are imposed, so every draw enters with *equal weight*.