mpshock

Cross-country monetary policy shocks in R

Charles Coverdale

London, United Kingdom

April 2026

The question

How do we make identified monetary policy shock series reproducible, comparable, and cross-country in a single R call?

Thirteen series, three countries, five identification strategies, one tidy schema.

Why it matters

- **Empirical macroeconomists** run impulse responses under several identification strategies as a standard robustness check
- **Policy institutions** (central banks, Treasuries) want cross-country comparability of monetary stance measures
- **Graduate students and replicators** lose hours assembling shock files before the interesting analysis begins
- **Forecasters and market strategists** feed shock series into nowcasting pipelines and scenario models

What is already out there

- **Personal replication files**: Nakamura-Steinsson on Harvard Dataverse, Bauer-Swanson on the SF Fed, Jarocinski-Karadi on personal pages
- **`hfdshocks`** (GitHub): single-country, unmaintained, no provenance metadata
- **Bank of England UKMPD spreadsheet**: live-maintained but UK-only, Excel format
- **Individual central-bank bulletins**: Atlanta Fed (Wu-Xia), RBA (Hambur-Haque, Beckers)

The gap: **no R package bundles shock series for more than one country in a common schema**.[1]

[1] Aeberhardt et al. (2024) report pairwise correlations of 0.4 to 0.7 across alternative US shock measures on common samples; reporting several identifications in parallel is now standard.

What mpshock offers

1. **Coverage**: 13 series across US (8), UK (3), and Australia (2), spanning narrative, high-frequency, sign-restriction, informationally-robust, and shadow-rate identifications
2. **Interface**: every series loads via `mp_shock(series, start, end)` and returns a tidy `data.frame` with the `mp_shock` S3 class; outputs compose via `rbind()`
3. **Provenance**: each series carries DOI, source URL, licence, and download date; pure data + computation, no runtime network, 240 unit tests

On CRAN since April 2026. Plugs directly into `lpirfs`, BVAR, and `vars`[2].

[2] Adammer (2019) *lpirfs*, Kuschnig & Vashold (2021) *BVAR*, Pfaff (2008) *vars*, all R packages for impulse response estimation.

US shocks: cumulative paths across four identifications

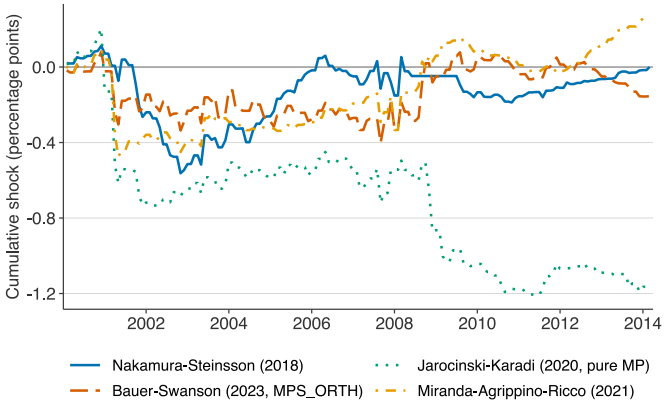

Cumulative monthly US monetary policy shocks under four identifications over a common sample. Exports: `mp_shock("nakamura_steinsson")`, `mp_shock("bauer_swanson")`, `mp_shock("jarocinski_karadi_mp")`.

US shock correlations: 0.4 to 0.7 across identifications

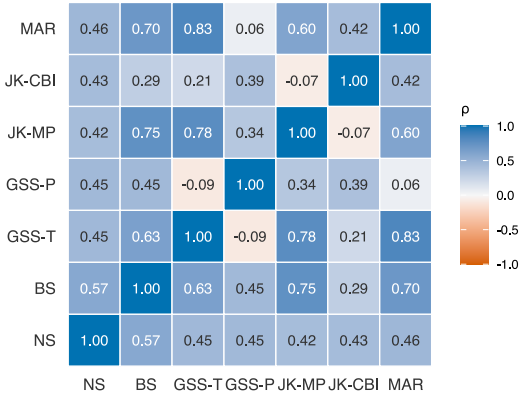

Pairwise Pearson correlations among seven US shock and stance series, common monthly sample 2000-02 to 2014-03. Exports: `mp_shock()`, `mp_align()`.

Coverage: thirteen series on one timeline

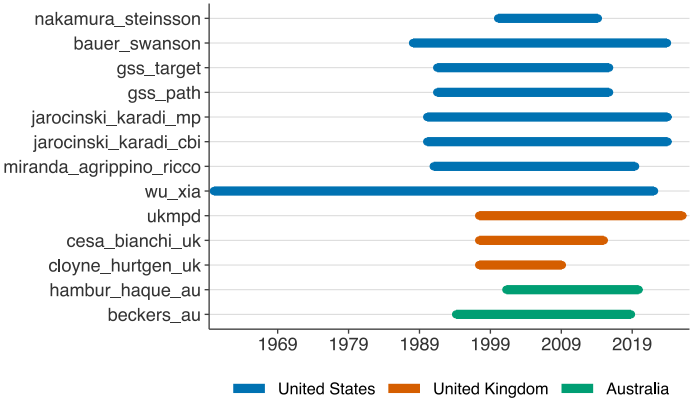

Time coverage of all thirteen bundled series across the US (8), UK (3), and Australia (2); each row marks the native publication window. Exports: `mp_list()`, `mp_source()`, `mp_shock()`.

Distributions: shock magnitudes by series

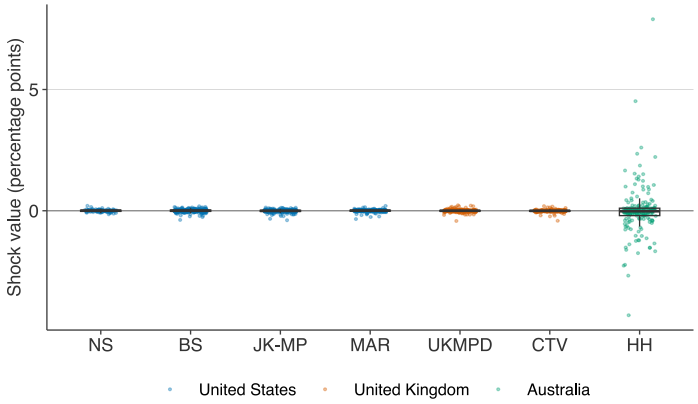

Distribution of nonzero monthly shock magnitudes by series, in native units. Narrative series show heavier right tails; high-frequency surprises are tighter. Exports: `mp_shock()`, `mp_cumulate()`.

Central formulas

High-frequency surprise (30-minute window around FOMC):

$$\mathrm{HFS}_t = f_t^{+} - f_t^{-} \tag{1}$$

Orthogonalised shock (Bauer-Swanson residual):

$$\mathrm{MPS\_ORTH}_t = \mathrm{MPS}_t - X_t' \hat{\beta} \tag{2}$$

Informationally-robust shock (Miranda-Agrippino-Ricco):

$$u_t = \mathrm{FF4}_t - \sum_{h=0}^{4} \gamma_h \, \Delta \mathbb{E}_t^{\mathrm{GB}}[y_{t+h}] \tag{3}$$

Package at a glance

Bundled series by country:

- **United States**: 8 series (Nakamura-Steinsson, Bauer-Swanson, GSS target, GSS path, Jarocinski-Karadi MP and CBI, Miranda-Agrippino-Ricco, Wu-Xia)
- **United Kingdom**: 3 series (UKMPD three-factor, Cesa-Bianchi, Cloyne-Hurtgen)
- **Australia**: 2 series (Hambur-Haque three-factor, Beckers)
- **Loader + metadata**: `mp_shock()`, `mp_list()`, `mp_source()`
- **Helpers**: `mp_align()`, `mp_to_quarterly()`, `mp_cumulate()`

Deps: `cli`, `stats`, `utils`. R $\geq$ 4.1.0.

Uniform output

Every series returns:
`date` | `shock` | `series`
Multi-factor series add
`path`, `qe`, or `term_premium`.
Outputs compose via `rbind()` into panel data.

Minimal working example

```r
library(mpshock)

# Three US identifications on a common monthly sample
ns  <- mp_shock("nakamura_steinsson", start = "2000-02-01")
bs  <- mp_shock("bauer_swanson",      start = "2000-02-01")
jk  <- mp_shock("jarocinski_karadi_mp", start = "2000-02-01")

# Compose into a panel for plotting or modelling
panel <- rbind(ns, bs, jk)

# Hand off to downstream IRF estimation
lpirfs::lp_lin_iv(endog_data, shock = panel$shock)
```

All three calls share the same three-column schema, so `rbind()` works without reshaping or renaming.

Cross-country monetary policy shocks, 2005 to 2022

Data. `bauer_swanson` (US MPS_ORTH), `ukmpd` Target factor (UK), `hambur_haque_au` action factor (AU). Each standardised by full-sample standard deviation of nonzero shocks before cumulating.

Question. *Do identified monetary policy shocks in the US, UK, and Australia move together around major global events, or does country-specific identification dominate?*

Why this case. First published cross-country cumulative shock chart from a single R pipeline[3]. Covers GFC, 2013 taper tantrum, COVID-19, and 2022-23 tightening cycle. Standardisation makes scales comparable despite different native units.

[3] No prior R package bundles shock series for more than one country. See paper section on cross-country shocks for full discussion.

Hero result: US, UK, and AU shocks move together through global shocks

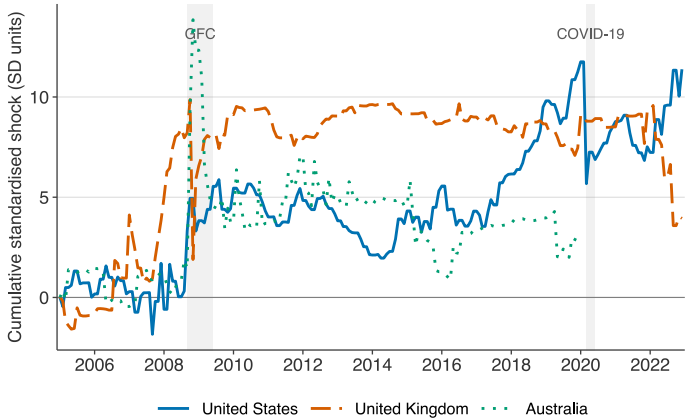

Cumulative standardised monetary policy shocks for the US, UK, and Australia. Shaded bands: GFC (Sep 2008 to Jun 2009) and COVID-19 (Mar to Jun 2020). Sources: Bauer-Swanson (US), UKMPD Target (UK), Hambur-Haque

What mpshock does not yet do

- **Three countries only**: Euro-area (Jarocinski-Karadi ECB, Altavilla EA-MPD) is the v0.2.0 priority
- **Monthly aggregation sums event surprises**: Bu, Rogers, and Wu (2021) argue for NA-coding in LP-IV estimation
- **No new identification**: mpshock is a data-curation project, not a new econometric method
- **Licences vary**: users cite the underlying paper for any series used in academic work

v0.2.0 roadmap: Euro-area shocks (Altavilla EA-MPD, JK ECB), Romer-Romer US narrative series, Bolhuis-Das-Yao cross-country panel subject to licence.

Contact, code, paper

Charles Coverdale

charles.f.coverdale@gmail.com
github.com/charlescoverdale/mpshock
charlescoverdale.github.io/publications

Acknowledgements. I thank the authors of each bundled series for making their data publicly available: Nakamura, Steinsson, Bauer, Swanson, Gürkaynak, Sack, Jarocinski, Karadi, Miranda-Agrippino, Ricco, Wu, Xia, Braun, Saha, Cesa-Bianchi, Thwaites, Vicondoa, Cloyne, Hürtgen, Hambur, Haque, and Beckers.

Scan for paper PDF