inflationkit

Inflation decomposition, persistence, and Phillips curves in R

Charles Coverdale

London, United Kingdom

May 2026

The question

How does an applied economist run the standard inflation diagnostic playbook in one R package?

Eleven functions, one `ik_` prefix, every dataset in any country.

Why it matters

- **Central banks** (Fed, ECB, BoE) report core, trim, persistence and Phillips curve estimates in every quarterly briefing
- **Sovereign treasuries and OBR-equivalents** track breakeven inflation and forecast bias for fiscal projections
- **Research departments** (NBER, CEPR, IMF) need replicable persistence and slope estimates across countries
- **Financial press and policy analysts** decompose headline spikes into supply versus demand each release day

What is already out there

- **`Inflation`** (CRAN): four core measures only, no decomposition, persistence, Phillips curves, or forecast evaluation; unmaintained since 2017[1]
- **Spreadsheet workbooks**: every applied team has its own; no version control, no testing, no citation
- **Bespoke scripts**: persistence and Phillips curve code is rewritten in every paper
- **Proprietary platforms** (Haver, Bloomberg): closed source, no transparency, no reproducibility

The gap: **no single CRAN package covers the standard inflation workflow end to end**.

[1] Iturbide & Tovar (2017), *Inflation: Core Inflation*, R package v1.4.

What inflationkit offers

1. **Coverage**: 11 functions across decomposition, core, persistence, Phillips curves, trend extraction, breakeven, and forecast evaluation
2. **Interface**: every function accepts arbitrary column names and returns an S3 object with a `print()` method; `rbind()` composes panels
3. **Provenance**: pure R, no compiled code, no API calls, no bundled data; over 200 tests; CRAN since March 2026

Methods follow @bryan1994measuring (core), @pivetta2007persistence (persistence), @stock2007why (trend), @mincer1969evaluation (forecast eval)[2].

[2] Bryan & Cecchetti (1994), *Measuring core inflation*, NBER WP 4303.

Decomposition: what is driving headline inflation

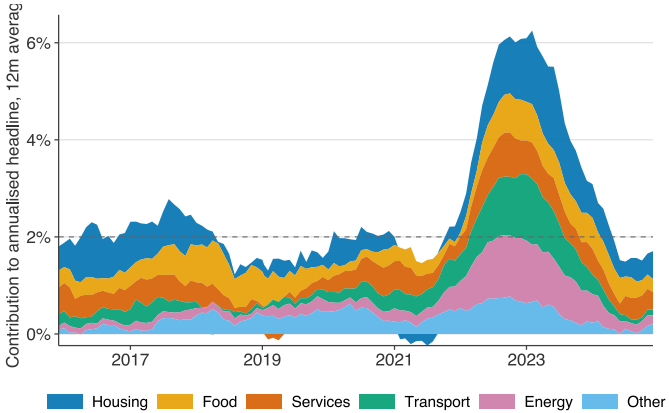

Component contributions sum to headline. Energy and transport drove the 2021-2022 peak; housing contributes slowly throughout. Exports: `ik_decompose()`, `ik_diffusion()`.

Core measures: four ways to strip the noise

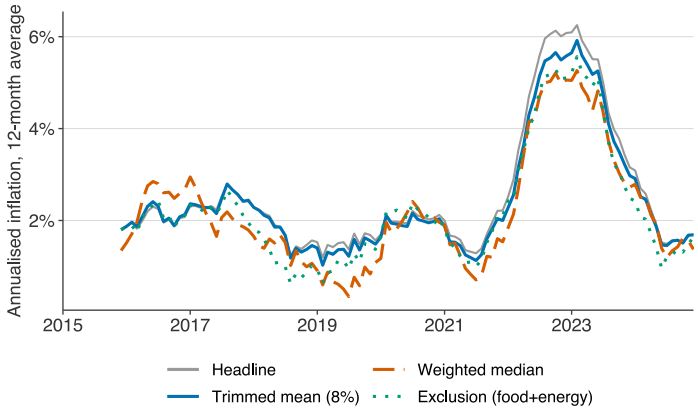

Trimmed mean, weighted median, exclusion, asymmetric trim. Distribution-based measures attenuate the spike more than fixed exclusion. Exports: `ik_core()`, `ik_sticky_flexible()`.

Phillips curves: slack, expectations, and the slope

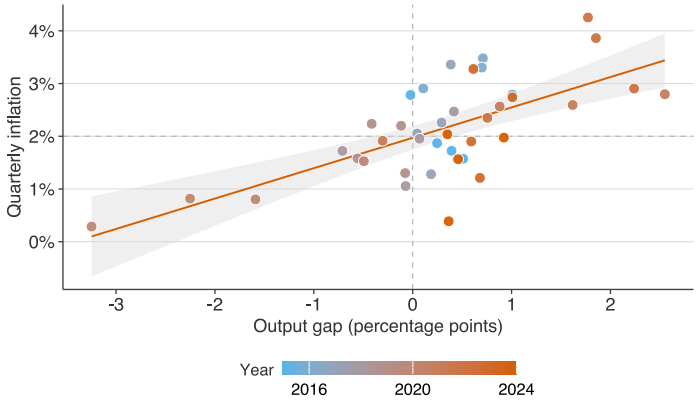

Traditional, expectations-augmented, and hybrid New Keynesian forms with HAC standard errors. Exports: `ik_phillips()`, `ik_persistence()`, `ik_breakeven()`.

Central formulas

Persistence (sum of AR coefficients):

$$\rho = \sum_{j=1}^{p} \phi_j, \qquad \pi_t = c + \sum_{j=1}^{p} \phi_j \pi_{t-j} + \varepsilon_t \tag{1}$$

Phillips curve (hybrid New Keynesian):

$$\pi_t = \alpha + \beta_f \, \mathbb{E}_t \pi_{t+1} + \beta_b \, \pi_{t-1} + \gamma \, x_t + \varepsilon_t \tag{2}$$

Mincer-Zarnowitz (forecast bias):

$$\pi_t = \alpha + \beta \, \hat{\pi}_t + \varepsilon_t, \qquad H_0 : \; (\alpha, \beta) = (0, 1) \tag{3}$$

Package at a glance

Function families:

- **Decomposition**: `ik_decompose`, `ik_core`, `ik_sticky_flexible`, `ik_diffusion`
- **Dynamics**: `ik_persistence`, `ik_phillips`
- **Trend**: `ik_trend` (HP, BN, exponential, MA)
- **Breakeven**: `ik_breakeven`
- **Forecast eval**: `ik_forecast_eval` (MZ, Nordhaus, DM)
- **Helpers**: `ik_compare`, `ik_sample_data`

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

Flexible column names

Every function accepts:
`date_col` | `item_col` | `weight_col` | `price_col`
ONS, Eurostat, BLS, ABS data flows in without reshaping.

Minimal working example

```
library(inflationkit)

# Five-line canonical workflow
d <- ik_sample_data("components")

ik_decompose(d)                                          # contributions
ik_core(d, method = "trimmed_mean", trim = 0.08)         # core
ik_persistence(headline_series)                          # half-life
ik_phillips(inflation = infl, slack = output_gap,        # PC slope
            type = "traditional", robust_se = "HAC")
```

Each call returns an S3 object with a `print()` method. Outputs compose via `rbind()` for panel-style analysis across countries or vintages.

Case study: the pandemic-era playbook on US CPI

Data. FRED series `CPIAUCSL` (headline) and `CPILFESL` (less food and energy), monthly, 2015 to 2026.

Question. *Was the 2021-2023 spike a relative-price shock or a regime shift in trend inflation?*

Why this case. Largest test of inflation targeting in three decades. Every central bank ran exactly this diagnostic sequence in real time[3].

[3] Reis (2022), *The burst of high inflation*, NBER WP 30616. Shapiro (2022), *Decomposing supply and demand driven inflation*, FRBSF WP.

US CPI, 2015 to 2026: headline, core, and HP trend

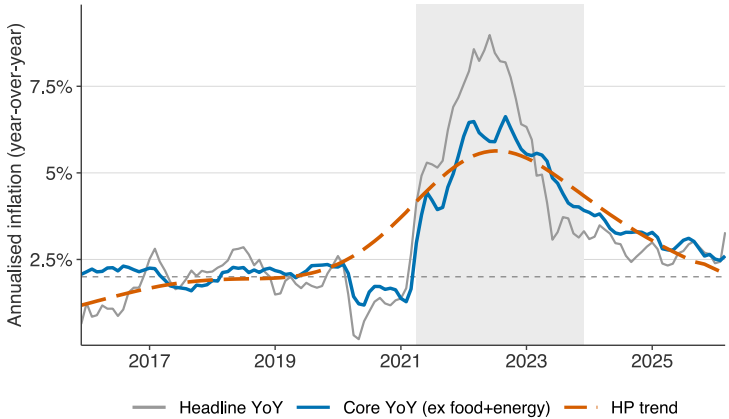

Headline (grey) peaks above 9 per cent in June 2022; core (blue) peaks at 6.7 per cent. HP trend (red dashed) captures the regime shift. `ik_persistence()` *on year-over-year headline returns sum-of-AR = 0.97. Source:*

Mincer-Zarnowitz: a naive AR(1) forecast was biased through the spike

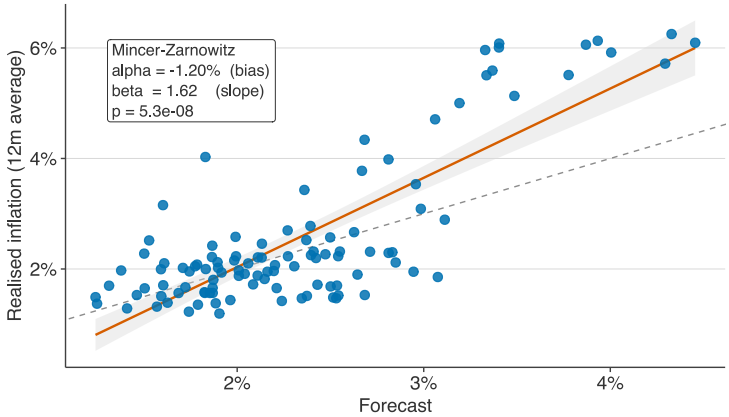

Realised inflation against an AR(1)-with-drift forecast, 120 months synthetic data. Slope below unity, intercept off zero: under-reaction and bias. `ik_forecast_eval()` *returns the test statistic and joint* p*-value.*

What inflationkit does not yet do

- **Single-equation OLS only**: no IV, no GMM, no @hazell2022slope cross-sectional design
- **No bundled price data**: pure computation; data still comes from ONS, Eurostat, BLS, ABS
- **No regional decomposition**: the package operates at the country level
- **No DSGE-style structural inference**: that is the job of `dynare4r` or bespoke Stan code

v0.2.0 roadmap: regional Phillips curve panel estimators (HSL-style cross-sectional design), expanded breakeven options across maturities, sticky-flexible helpers for the full Atlanta Fed component map.

Contact, code, paper

Charles Coverdale

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

Acknowledgements. CPI data: Federal Reserve Bank of St. Louis (FRED, public). Methodological debts to Bryan, Cecchetti, Pivetta, Stock, Watson, Mincer, Zarnowitz, Diebold, Mariano, Hodrick, Prescott, and the Cleveland Fed trim-CPI team. Thanks to CRAN maintainers and reviewers of *inflationkit* v0.1.0.

Scan for paper PDF