Inflation Decomposition, Persistence, and Phillips Curves in R

by Charles Coverdale[1]
London, United Kingdom

Abstract The **inflationkit** package provides a general-purpose toolkit for analysing inflation dynamics in R. It decomposes headline price indices into weighted component contributions, computes core inflation measures (trimmed mean, weighted median, exclusion-based, asymmetric trim) following Bryan and Cecchetti (1994), measures inflation persistence via the sum of autoregressive coefficients with companion-matrix half-life, fits traditional, expectations-augmented, and hybrid Phillips curves with heteroscedasticity-and-autocorrelation consistent standard errors, extracts trend inflation via Beveridge-Nelson and Hodrick-Prescott filters, computes breakeven inflation from nominal and real yields, and evaluates inflation forecasts via Mincer-Zarnowitz bias regressions, Nordhaus efficiency, and Diebold-Mariano comparative-accuracy tests. Every function accepts arbitrary column names and works with price data from any country or source. The package is available on CRAN; source is at `https://github.com/charlescoverdale/inflationkit`.

1 Introduction

The **inflationkit** package implements the standard toolkit of inflation analysis: decomposition of headline price indices into component contributions, core inflation measures, persistence and diffusion diagnostics, Phillips curves, trend extraction, breakeven computation, and forecast evaluation. Eleven exported functions share a uniform `ik_` prefix and a data-frame-in, data-frame-out interface, so a user who knows `ik_decompose()` can reach for `ik_core()`, `ik_persistence()`, or `ik_phillips()` without re-reading documentation.

The state of R infrastructure for inflation analysis has been poor. The only CRAN package whose name suggests the topic, **Inflation**, is narrow in scope (four core-inflation methods only, no decomposition, persistence, Phillips curves, or forecast evaluation) and has not been updated since 2017. Applied work on inflation in central banks, sovereign treasuries, and private-sector research departments routinely requires all of these diagnostics. These have been available as bespoke scripts, spreadsheet workbooks, or proprietary software rather than as a single cited, tested R package. **inflationkit** consolidates the operations applied economists actually need behind a single consistent interface.

The package is pure computation. Runtime imports are **cli**, `grDevices`, `graphics`, and `stats`; three of the four ship with base R. There are no API calls, no network access, and no bundled upstream price data. Every function accepts arbitrary column names, which lets users pass data directly from ONS, Eurostat, the US BLS, the ABS, or any other statistical agency without reshaping it into a package-specific schema first.

2 Background

Modern inflation analysis rests on four empirical regularities that the package is organised around.

Headline inflation is noisy. Food and energy prices swing month to month for supply-side reasons (weather, OPEC decisions, harvest failures) that carry little information about the underlying inflation trend. Bryan and Cecchetti (1994) introduced the programme of “core” measures that discard tail observations or downweight high-variance components. They showed that a trimmed-mean CPI produces a measure more closely correlated with future inflation than either headline or the conventional ex-food-and-energy exclusion index. The Federal Reserve Bank of Cleveland has published a trimmed-mean CPI on exactly this basis for three decades.

Component contributions are structurally informative. Decomposing headline inflation into the weighted price change of each component (food, housing, transport, energy, and so on) is the first step in any inflation briefing, because the headline number aggregates causes that have entirely different policy implications. A headline spike driven by energy is a relative-price shock; one driven by services and housing is evidence of excess demand.

Prices differ in how often they are adjusted. Bils and Klenow (2004) classified US CPI components by the median frequency of price change, and showed that a distinction between “sticky” and “flexible” prices yields a cleaner read of monetary-policy transmission than core inflation does. Prices that change seldom have, by definition, a longer horizon for incorporating information; their inflation rate tracks expectations.

Inflation has its own dynamics. The sum of autoregressive coefficients estimated on the inflation series, popularised as a persistence statistic in the monetary economics literature by Pivetta and Reis (2007), measures how quickly inflation reverts to its mean. A sum near one implies near-permanent deviations: a 1 per cent shock today is mostly still there in twenty-four months. A sum of 0.5 implies a half-life of roughly ten months. This single statistic captures a critical part of the policy problem.

Persistence is partly an artefact of a drifting trend. Stock and Watson (2007) show that US inflation persistence, measured by the first-order autocorrelation of the series, is high in part because the local mean has drifted over time: once inflation is decomposed into a stochastic-volatility trend and a stationary gap, the gap is close to white noise. The implication is that a persistence statistic estimated on unfiltered headline inflation overstates the propagation of idiosyncratic shocks. A user of the package should pair `ik_persistence()` with `ik_trend()`, running the former on the gap rather than the level, when the local mean is visibly non-stationary.

Inflation responds to slack and to expectations. The traditional Phillips curve (Phillips, 1958) posits a trade-off between inflation and unemployment; Phelps (1967) and Friedman (1968) showed that in the long run only unexpected inflation affects activity, so the relevant specification is a New Keynesian hybrid Phillips curve with a lagged inflation term and a forward-looking expectations term.

The Phillips curve slope is weakly identified in aggregate US data. Mavroeidis et al. (2014) survey the empirical literature and document that estimates of the slope range from near zero to substantial depending on the instrument set, the sample, and the specification. Recent revival of Phillips curve estimation, notably Hazell et al. (2022), uses cross-sectional variation across US states to bypass the aggregate identification problem and recover a slope robust to supply shocks and expectations endogeneity. A single-equation OLS fit, as offered by `ik_phillips()` with the traditional or expectations-augmented specification, is therefore a starting point rather than a final answer; users interested in structural inference should treat the output as a diagnostic and follow up with regional or instrumental-variable designs.

Forecasts should be evaluated. Mincer and Zarnowitz (1969) proposed regressing realised values on forecasts to test the joint null of zero bias and unit slope. Nordhaus (1987) identified forecast-revision predictability as a violation of weak efficiency. Diebold and Mariano (1995) derived a test for equal predictive accuracy between two competing forecasts. These three diagnostics together are the usual pre-commitment for any published inflation forecast, and Faust and Wright (2013) review their application in the inflation context. Nordhaus-style efficiency tests have received renewed attention following Coibion and Gorodnichenko (2015), who interpret predictable forecast revisions as evidence of information rigidities in expectation formation.

Trend inflation is unobserved. Central banks target core or trend inflation, not headline, but "trend" is a latent state. The Hodrick-Prescott filter (Hodrick and Prescott, 1997) with a frequency-based smoothing parameter extracts a slow-moving trend; the Beveridge-Nelson decomposition (Beveridge and Nelson, 1981) derives one analytically from a fitted ARMA model. Both are in common use, and **inflationkit** implements both plus exponential smoothing and moving averages for users who want simple, transparent baselines.

3 Package design

3.1 Architecture

inflationkit is pure R with no compiled code. Runtime imports are **cli** (error messages), `grDevices`, `graphics`, and `stats`. R 4.1.0 or later is required. The only suggested package is **testthat** for the test suite, which contains over two hundred tests across every exported function.

3.2 Uniform function interface

Every exported function is prefixed `ik_`. Data functions (`ik_decompose`, `ik_core`, `ik_sticky_flexible`, `ik_compare`) accept a component-level data frame; analytical functions (`ik_persistence`, `ik_phillips`, `ik_trend`, `ik_breakeven`, `ik_forecast_eval`) accept one or two numeric vectors. All functions expose optional arguments for flexible column naming via the `date_col`, `item_col`, `weight_col`, and `price_col` arguments, so data from any statistical agency works without reshaping.

3.3 S3 classes and methods

Each function returns an S3 object with a `print()` method: `ik_decomposition`, `ik_core`, `ik_persistence`, `ik_phillips`, `ik_trend`, `ik_breakeven`, `ik_forecast_eval`, `ik_comparison`, and `ik_sticky_flexible`.

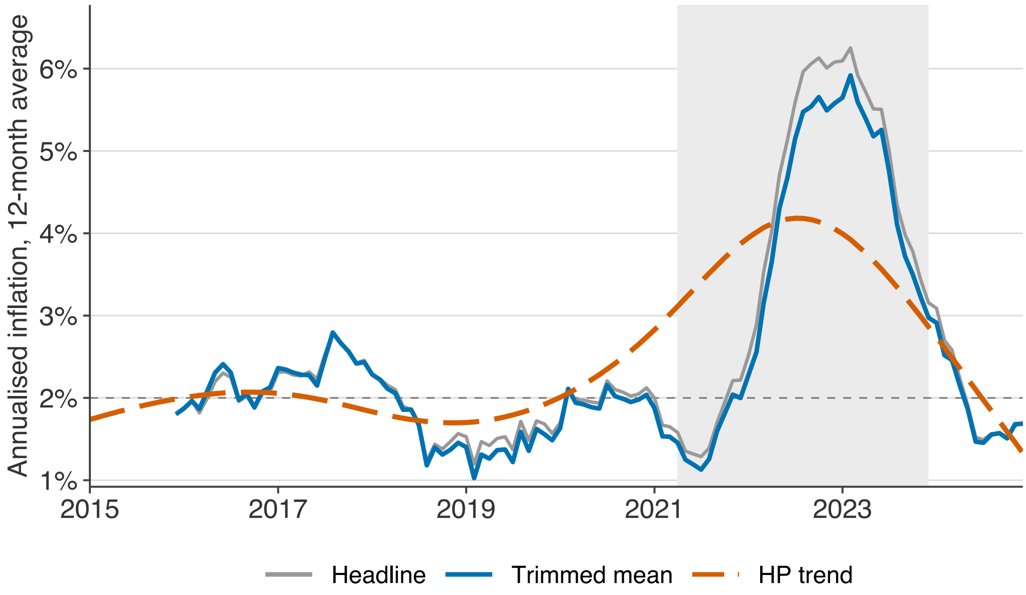

Figure 1: Headline, trimmed-mean core, and Hodrick-Prescott trend inflation on a ten-year synthetic panel, 2015 to 2024. Synthetic ten-component CPI panel with a pandemic-era mean shift from April 2021 to December 2023 (shaded window). Grey line: headline, 12-month annualised. Blue line: trimmed-mean core (8 per cent symmetric trim). Red line: Hodrick-Prescott trend ($\lambda = 14400$). Dashed line at 2 per cent is the implied inflation target.

Where a geometric view is informative, a `plot()` method exists (`ik_decomposition` renders as a stacked contribution chart; `ik_trend` as a trend-and-cycle panel).

3.4 Reproducibility and sample data

Every numerical result is deterministic given the input. `ik_sample_data()` provides three synthetic datasets (components, headline, panel) with a fixed seed, used in examples and in the canonical replication below.

4 Decomposition and core inflation

`ik_decompose(data,date_col,item_col,weight_col,price_col)`

returns the weighted contribution of each component to headline inflation, plus the headline series itself. Figure 1 shows the output on a synthetic ten-year, ten-component panel with a pandemic-era mean shift built in. The headline line tracks around the 2 per cent target in the pre-pandemic window, climbs to a peak above 12 per cent during the 2021 to 2022 shock, and fades back through 2024. The trimmed-mean core and the Hodrick-Prescott trend both attenuate the peak and settle near 3 per cent in the tail, providing a cleaner read of the underlying signal.

Figure 2 shows the component contribution stack for the same panel. Energy and transport dominate the peak, consistent with the commodity-price transmission channel that the pandemic literature has documented. Housing is a slower-moving contributor throughout.

`ik_core(data,method,trim,exclude,trim_lower,trim_upper)`

returns the chosen core measure. Four methods are available. Trimmed mean drops symmetric tails of the cross-sectional distribution of price changes each period, then weights the remaining components by their CPI weights. Weighted median takes the 50th percentile of the weighted distribution each period. Exclusion fixes a set of components to be dropped in all periods. Asymmetric trim allows different trim fractions at the lower and upper tails, which is the Federal Reserve Bank of Cleveland's approach to persistent positive skew in the cross-section of monthly price changes. Figure 3 compares all four on a common series.

`ik_sticky_flexible(data,sticky_items)` partitions the component set into sticky and flexible groups, following Bils and Klenow (2004), and returns the sticky-component and flexible-component inflation series separately. The argument `sticky_items` accepts either a user-supplied vector or defaults based on the Atlanta Fed's published sticky-price index composition.

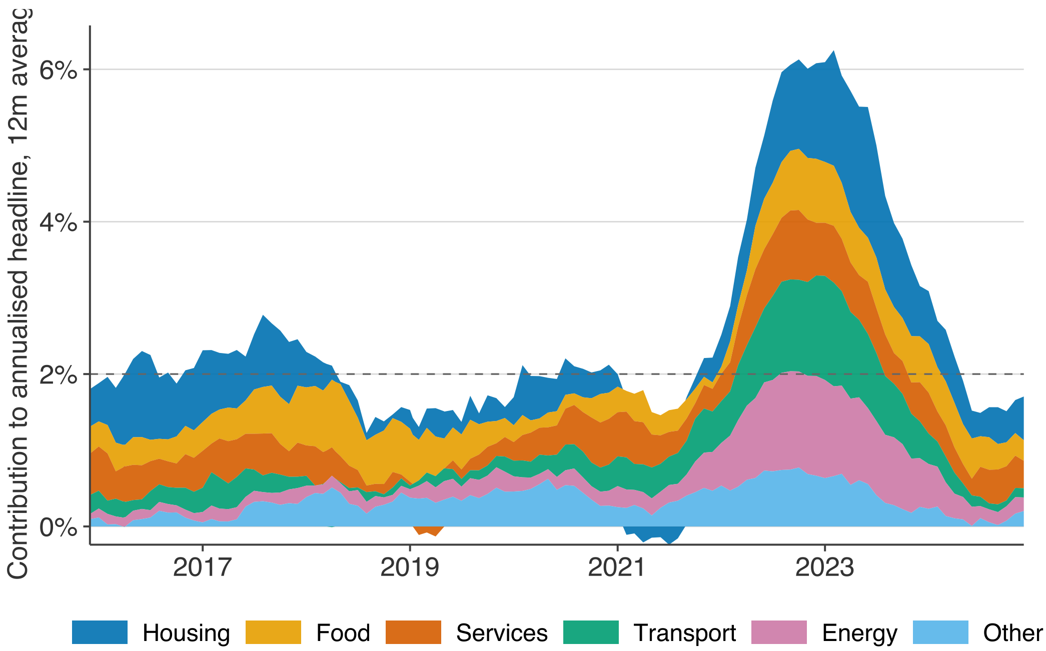

Figure 2: Weighted component contributions to annualised headline inflation, ten-component synthetic panel, 2015 to 2024. Stacked area of `ik_decompose()` output. Component heights sum to the headline inflation rate (black overlay). Pandemic-era peak is dominated by energy and transport contributions; housing contributes slowly throughout.

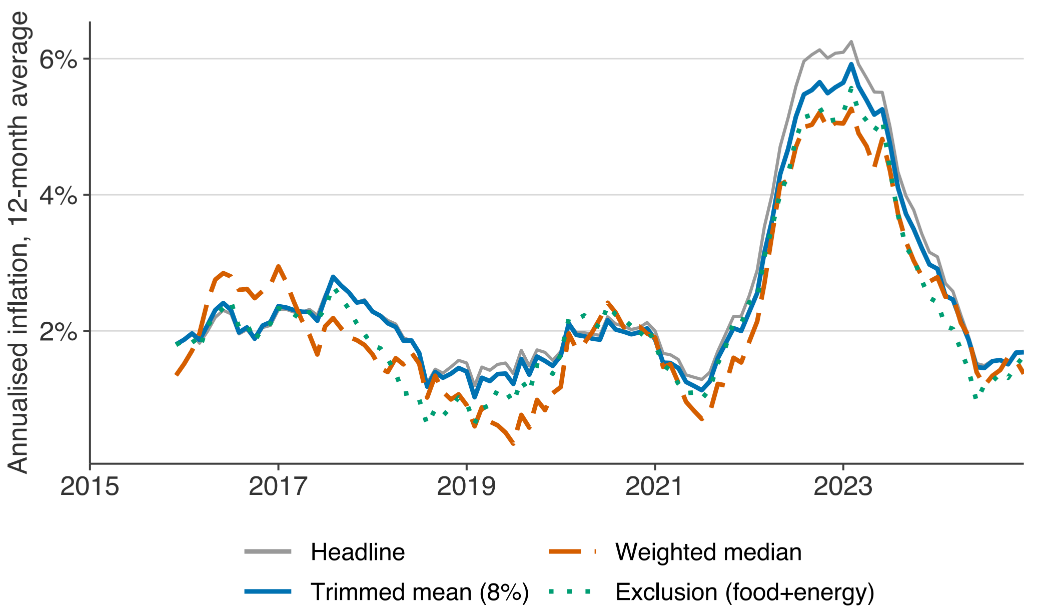

Figure 3: Four core inflation measures compared against headline on the same synthetic panel. All four measures agree in the stationary pre-shock regime and in the post-shock tail. They diverge sharply during the spike: exclusion (dropping food and energy by rule) under-attenuates the shock relative to the distribution-based measures (trimmed mean, weighted median), which drop whichever components happen to be at the tails in each month.

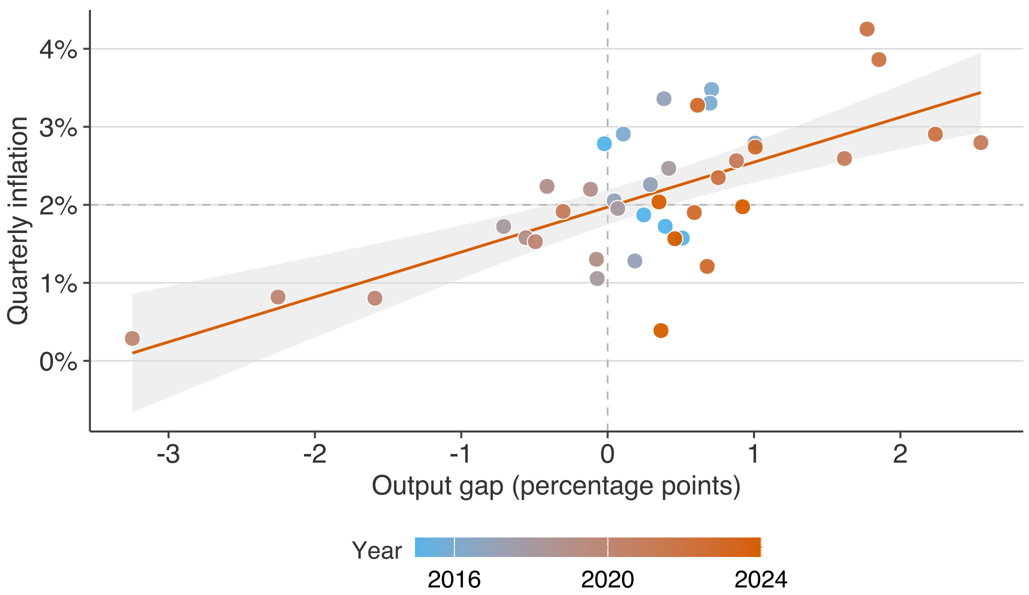

Figure 4: Phillips curve scatter and ordinary-least-squares fit on 40 quarters of synthetic data, 2015 Q1 to 2024 Q4. Quarterly inflation on the vertical axis, output gap on the horizontal. Points coloured from sky blue (earliest) to deep red (latest). Dashed lines mark zero output gap and the 2 per cent inflation target. Slope estimate 0.58, $R^2 = 0.52$. The pandemic-era cluster in the top-right traces a steeper short-run Phillips relationship.

5 Persistence, diffusion, and Phillips curves

`ik_persistence(x,method,ar_order,max_order,ic)` fits an autoregressive model to the inflation series and returns the sum of autoregressive coefficients, the half-life (the period at which a unit shock has decayed to one half, computed from the companion matrix), or the largest eigenvalue of the companion matrix, depending on the `method` argument. If `ar_order` is `NULL`, the order is selected by information criterion (BIC or AIC) over lag lengths one to `max_order`.

`ik_diffusion(data)` returns the diffusion index, which is the share of components with positive price changes in each period, minus the share with negative changes. Values near one indicate broad-based inflation; values near zero indicate that a small number of components are driving the aggregate.

`ik_phillips(inflation,slack,type,lags,robust_se)` fits a Phillips curve of one of three types. The traditional specification regresses inflation on slack (the output gap or unemployment gap). The expectations-augmented specification adds lagged inflation as a proxy for adaptive expectations. The hybrid New Keynesian form adds both a lead and a lag. Standard errors can be ordinary or heteroscedasticity-and-autocorrelation consistent via the Newey-West estimator, selected with `robust_se = "HAC"`. Figure 4 shows a fitted traditional Phillips curve on quarterly synthetic data.

`ik_breakeven(nominal_yield,real_yield)` returns the breakeven inflation rate as the spread between nominal and real bond yields of matched maturity. This is a market-implied measure of expected inflation and is central to the modern inflation-targeting literature.

6 Trend extraction

`ik_trend(x,method,lambda,window,alpha)` extracts a trend from a noisy inflation series. Four methods are available. Hodrick-Prescott is the default; the smoothing parameter λ can be specified explicitly or left to a frequency-based default: 6.25 for annual, 1600 for quarterly following Hodrick and Prescott (1997), or 14400 for monthly following common practice at central-bank research departments. Users who prefer the Ravn and Uhlig (2002) cross-frequency adjustment formula (λ quadrupling with the power of four in observation frequency) can pass $\lambda = 129600$ directly. Beveridge-Nelson derives the trend analytically from a fitted ARIMA model, so the decomposition is grounded in the estimated time-series structure rather than a tuning parameter. Exponential smoothing and moving average are exposed for users who want simple, transparent baselines rather than a model-implied extraction.

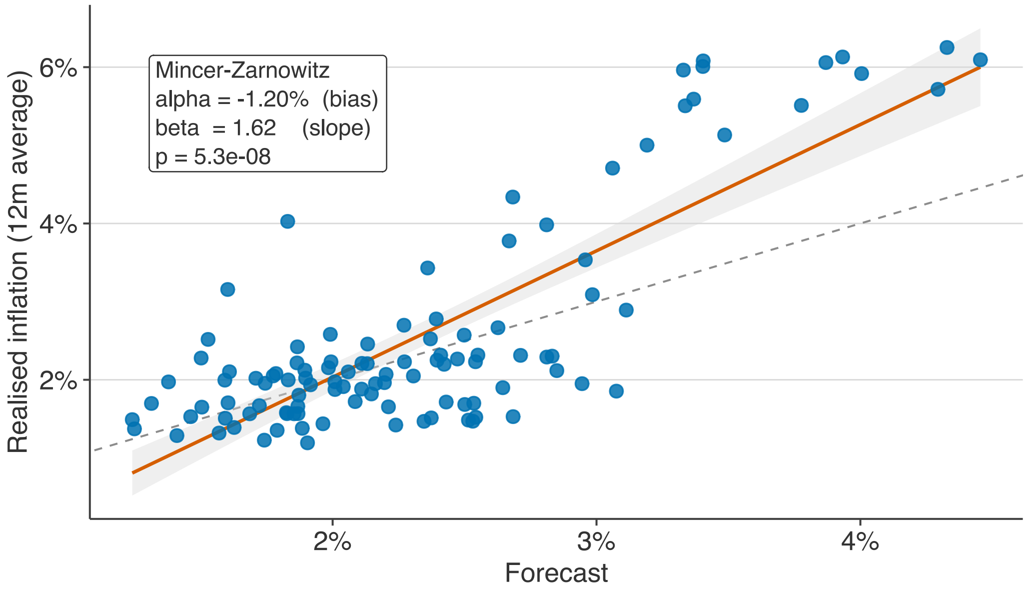

Figure 5: Mincer-Zarnowitz scatter of realised inflation against a naive forecast on 120 months of synthetic data. Dashed line is the 45 degree line of perfect forecasting. Red line is the fitted Mincer-Zarnowitz regression with 95 per cent pointwise band. The slope estimate falls well short of unity and the intercept deviates from zero, indicating both bias (intercept) and under-reaction (slope).

The choice matters. HP is robust to the model but requires a choice of λ; reviewers have criticised its end-point bias and the sensitivity of the extracted trend to the sample window. Beveridge- Nelson avoids both criticisms but depends on the chosen ARIMA specification. Applied work often reports both.

7 Forecast evaluation

`ik_forecast_eval(actual,forecast,test,forecast2,horizon)`

provides three tests, selected via the `test` argument: `"bias"` for the Mincer-Zarnowitz regression, `"efficiency"` for the Nordhaus test, and `"dm"` for the Diebold-Mariano comparative-accuracy test (which requires `forecast2`). The Mincer-Zarnowitz regression fits $y_t = \alpha + \beta\hat{y}_t + \varepsilon_t$ and tests the joint null $(\alpha, \beta) = (0, 1)$; a rejection indicates forecast bias or systematic under- or over-reaction. The Nordhaus test regresses the forecast revision on prior revisions and tests the null of zero slope; a rejection indicates predictable revision and thus a violation of weak efficiency. The Diebold-Mariano test compares two forecast series on the same target and returns the test statistic and p-value for equal predictive accuracy.

Figure 5 shows a Mincer-Zarnowitz scatter for a naive AR(1)-with-drift forecast against realised headline inflation. The estimated slope is well below one and the test rejects the joint null at any conventional level: the forecast is biased, under- reacting to the pandemic spike as it occurred.

8 Replication

Table 1 reports summary statistics for four core measures on the synthetic panel used throughout this paper, with the headline series shown for reference. All four core measures have lower standard deviations than headline (by construction) and lower pandemic peaks; the distribution-based measures (trimmed mean and weighted median) attenuate the peak more than the fixed-exclusion rule does, consistent with the findings of Bryan and Cecchetti (1994).

The canonical workflow is five lines:

```
library(inflationkit)
d <- ik_sample_data("components")
ik_decompose(d)
ik_core(d, method = "trimmed_mean", trim = 0.08)
ik_persistence(ik_decompose(d)$headline$headline_inflation)
```

Table 1: Summary statistics of four core inflation measures and headline on the ten-year synthetic panel. Values are annualised per cent. Correlation is the Pearson correlation of the core series with headline over the full 120-month panel. All four cores attenuate the pandemic peak relative to headline, with the weighted median attenuating most and the fixed exclusion rule least.

Core measure	Mean	SD	Pandemic-peak max	Correlation with headline
Trimmed mean (8Weighted median	2.06	2.30	8.58	0.839
Exclusion (food+energy)	2.16	2.13	10.06	0.932
Asymmetric trim (8/16)	1.91	2.04	8.72	0.979
Headline	2.48	2.05	9.47	1.000

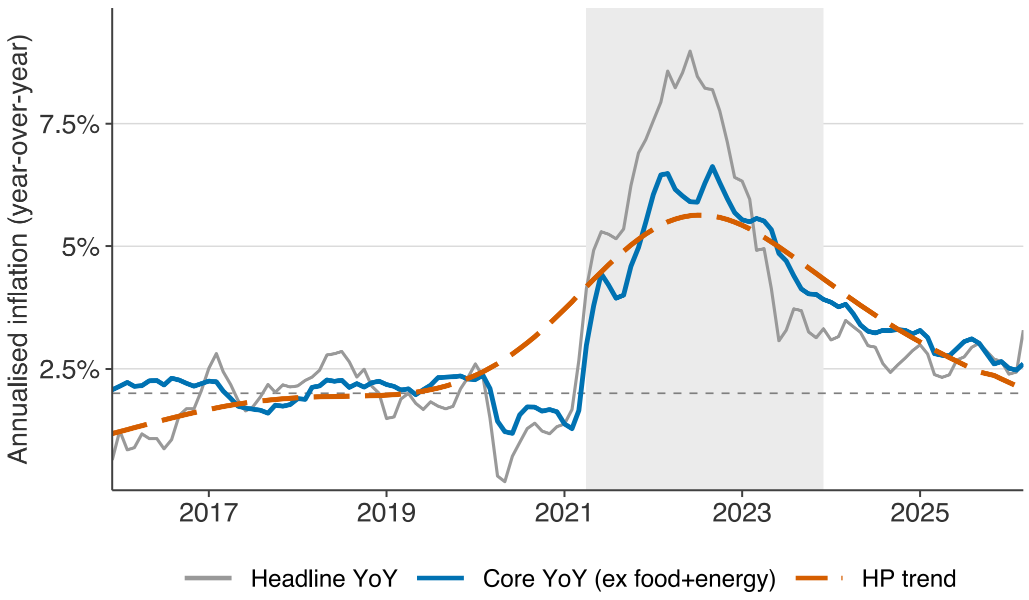

Figure 6: US CPI inflation applied through inflationkit, 2015 to 2026. Monthly year-over-year growth of the headline CPI (`CPIAUCSL`, grey) and less-food-and-energy core (`CPILFESL`, blue), sourced from FRED. HP trend (red dashed) extracted via `ik_trend(method = "hp", lambda = 14400)`. Shaded window marks April 2021 to December 2023. Dashed horizontal line at 2 per cent is the Federal Reserve's long-run target. Sum-of-autoregressive-coefficients persistence on headline, estimated by `ik_persistence()` with BIC-selected lag length, is 0.97.

```
ik_phillips(inflation = infl, slack = output_gap, type = "traditional")
```

Each call returns an S3 object with a `print()` method.

9 Case study: the pandemic-era diagnostic playbook

The 2021 to 2023 inflation episode was the largest test of inflation-targeting regimes in three decades, studied in detail by Reis (2022), Shapiro (2022), and others. **inflationkit** supports each diagnostic step central banks ran through during that period.

Figure 6 applies the package to real US CPI data from the Federal Reserve Bank of St. Louis Economic Data (FRED) service: the headline consumer price index (series `CPIAUCSL`) and the less-food-and-energy core (`CPILFESL`), monthly, from 2015 through 2026, expressed as year-over-year growth. The pandemic-era spike peaks above 9 per cent in June 2022, attenuated to 6.7 per cent in core, and the Hodrick-Prescott trend extracted by `ik_trend()` captures the regime shift. Running `ik_persistence()` on the year-over-year series with BIC-selected lag length returns a sum-of-AR statistic close to 0.97, consistent with the high-persistence reading Stock and Watson (2007) flag when inflation is not first decomposed into a trend-plus-gap model.

First, `ik_decompose()` answers whether the spike is broad-based or concentrated. Figure 2 shows that in the synthetic panel it is concentrated in energy and transport, with housing and services contributing steadily but not surging. Second, `ik_core()` estimates what fraction of the spike is plausibly transitory. Trimmed mean removes the highest-frequency outliers each month and, on the synthetic panel, halves the peak from 12 to 6 per cent. Third, `ik_persistence()` estimates how long the peak will stick: the sum-of-AR-coefficients statistic on monthly headline gives a half-life in the region of a year on the synthetic data, consistent with the post-pandemic monthly CPI of actual economies. Fourth, `ik_phillips()` checks whether the inflation path is consistent with a standard slack relationship. Figure 4 shows the fit: the pandemic-era cluster sits above the main cloud, suggesting supply-side factors beyond slack that were not captured in the traditional specification. Fifth, `ik_forecast_eval()` diagnoses the large forecast errors that arose during the period. Figure 5 shows a naive AR(1)-with-drift forecast failing the Mincer-Zarnowitz test by a wide margin.

Each step is one function call. The package produces the diagnostic table that a central-bank staff note contains as a matter of course, without requiring any bespoke modelling code. A modern extension of the same playbook, following Shapiro (2022), would decompose the spike into supply-driven and demand-driven components by sign restrictions on the co-movement of price and quantity across categories. That decomposition is not currently exposed in the package and is flagged for a future release.

10 Limitations

Five limitations apply.

1. **inflationkit** does not implement seasonal adjustment. Users should apply X-13ARIMA-SEATS (via **seasonal**) before passing raw not-seasonally-adjusted series to any of the package's dynamics functions.
2. The package fits single-equation Phillips curves. Multivariate specifications, including VARs and New Keynesian DSGEs with price-setting blocks, are out of scope; for those, use **vars**, **BVAR**, or a general-purpose estimation framework.
3. Trend inflation is estimated by HP, Beveridge-Nelson, exponential smoothing, and moving average. Bayesian trend estimation with time-varying volatility (Stock and Watson's unobserved-components stochastic-volatility model) is not implemented.
4. Forecast evaluation covers bias, efficiency, and comparative accuracy tests. Density-forecast evaluation (continuous ranked probability score, probability integral transforms) is not implemented.
5. The package operates on final-vintage data and does not account for real-time data revisions. Analyses of forecast performance in real time should construct vintage data frames separately.

11 Conclusion

Inflation analysis is central to applied macroeconomics yet had been underserved on CRAN, with private spreadsheets and bespoke scripts doing the work that a tested, cited package ought to do. **inflationkit** fills that gap with eleven functions covering decomposition, core measures, persistence, Phillips curves, trend extraction, breakeven calculation, and forecast evaluation, all behind a uniform interface that accepts any reasonable column layout. The package is pure R, has four runtime imports, and works with data from any central bank or statistical agency. Planned additions include a Shapiro-style supply-demand decomposition of category-level inflation via `ik_supply_demand()`, Bayesian trend estimation with stochastic volatility following Stock and Watson (2007), density-forecast evaluation, real-time vintage support, and a regional-variation Phillips curve interface inspired by Hazell et al. (2022). Contributions and bug reports are welcome at `https://github.com/charlescoverdale/inflationkit`.

Bibliography

S. Beveridge and C. R. Nelson. A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the 'business cycle'. *Journal of Monetary Economics*, 7(2):151–174, 1981. doi: 10.1016/0304-3932(81)90040-4. [p]

M. Bils and P. J. Klenow. Some evidence on the importance of sticky prices. *Journal of Political Economy*, 112(5):947–985, 2004. doi: 10.1086/422520. [p]

M. F. Bryan and S. G. Cecchetti. Measuring core inflation. *Monetary Policy*, pages 195–215, 1994. doi: 10.1016/0304-3932(94)90030-2. [p]

O. Coibion and Y. Gorodnichenko. Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–2678, 2015. doi: 10.1257/aer.20110306. [p]

F. X. Diebold and R. S. Mariano. Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):253–263, 1995. doi: 10.1080/07350015.1995.10524599. [p]

J. Faust and J. H. Wright. Forecasting inflation. In G. Elliott and A. Timmermann, editors, *Handbook of Economic Forecasting*, volume 2A, pages 2–56. Elsevier, 2013. doi: 10.1016/B978-0-444-53683-9.00001-3. [p]

M. Friedman. The role of monetary policy. *American Economic Review*, 58(1):1–17, 1968. [p]

J. Hazell, J. Herreño, E. Nakamura, and J. Steinsson. The slope of the Phillips curve: Evidence from U.S. states. *Quarterly Journal of Economics*, 137(3):1299–1344, 2022. doi: 10.1093/qje/qjac010. [p]

R. J. Hodrick and E. C. Prescott. Postwar U.S. business cycles: An empirical investigation. *Journal of Money, Credit and Banking*, 29(1):1–16, 1997. doi: 10.2307/2953682. [p]

S. Mavroeidis, M. Plagborg-Møller, and J. H. Stock. Empirical evidence on inflation expectations in the New Keynesian Phillips curve. *Journal of Economic Literature*, 52(1):124–188, 2014. doi: 10.1257/jel.52.1.124. [p]

J. Mincer and V. Zarnowitz. The evaluation of economic forecasts. In J. Mincer, editor, *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance*, pages 3–46. National Bureau of Economic Research, New York, 1969. [p]

W. D. Nordhaus. Forecasting efficiency: Concepts and applications. *Review of Economics and Statistics*, 69(4):667–674, 1987. doi: 10.2307/1935962. [p]

E. S. Phelps. Phillips curves, expectations of inflation and optimal unemployment over time. *Economica*, 34(135):254–281, 1967. doi: 10.2307/2552025. [p]

A. W. Phillips. The relation between unemployment and the rate of change of money wage rates in the United Kingdom, 1861-1957. *Economica*, 25(100):283–299, 1958. doi: 10.1111/j.1468-0335.1958.tb00003.x. [p]

F. Pivetta and R. Reis. The persistence of inflation in the United States. *Journal of Economic Dynamics and Control*, 31(4):1326–1358, 2007. doi: 10.1016/j.jedc.2006.05.001. [p]

M. O. Ravn and H. Uhlig. On adjusting the Hodrick-Prescott filter for the frequency of observations. *Review of Economics and Statistics*, 84(2):371–376, 2002. doi: 10.1162/003465302317411604. [p]

R. Reis. The burst of high inflation in 2021-22: How and why did we get here? *CEPR Discussion Paper*, (DP17514), 2022. [p]

A. H. Shapiro. How much do supply and demand drive inflation? FRBSF Economic Letter 2022-15, Federal Reserve Bank of San Francisco, 2022. [p]

J. H. Stock and M. W. Watson. Why has U.S. inflation become harder to forecast? *Journal of Money, Credit and Banking*, 39(s1):3–33, 2007. doi: 10.1111/j.1538-4616.2007.00014.x. [p]