Nowcast: Economic Nowcasting with Bridge Equations and Real-Time Evaluation in R

by Charles Coverdale

Abstract The nowcast package provides bridge-equation nowcasting in R. It estimates current-quarter values of a low-frequency macroeconomic variable (typically GDP) from higher-frequency indicators (retail sales, industrial production, payrolls, sentiment) that are released weeks or months ahead of the official target. Mixed-frequency alignment handles the ragged edge of real-time data. Pseudo-real-time backtesting over expanding or rolling windows produces an out-of-sample performance record, and the Diebold-Mariano test with the Harvey-Leybourne-Newbold finite-sample correction compares competing specifications on the same target. The package is pure R with cli as the only non-base import, is available on CRAN, and plugs into any source of quarterly or monthly data.

1 Introduction

The **nowcast** package implements the standard workflow for bridge-equation nowcasting: mixed-frequency alignment, estimation with optional autoregressive terms, and pseudo-real-time evaluation. Ten exported functions share a uniform `nc_` prefix and accept plain data frames of dates and values, with no hard tie to **xts**, **zoo**, or any data-vendor schema.

Nowcasting (the practice of estimating a macroeconomic variable before it is officially released) has been standard in central banks since the early 2000s. GDP is typically published six to eight weeks after the quarter closes, but monthly indicators relevant to GDP arrive much sooner: US retail sales with a two- to three-week lag, industrial production with a two-week lag, nonfarm payrolls in the first week of the following month. The appeal of a nowcast is that it converts the flow of monthly releases during a quarter into a continuously updated estimate of the target.

R has had partial support for this workflow. **bridgr** implements bridge equations within the tidyverse ecosystem. **midasr** covers the related MIDAS family of mixed-frequency models. **bigtime** fits large-scale time-series models including sparse VARs at scale. None of these covers the full evaluation pipeline that a nowcasting analyst needs: pseudo-real-time backtesting with expanding or rolling windows, and pairwise Diebold-Mariano testing with the Harvey-Leybourne-Newbold finite-sample correction for small samples. **nowcast** fills that gap.

2 Background

The mixed-frequency problem. A quarterly target (GDP) must be aligned with monthly indicators. The standard choice, following Mariano and Murasawa (2003), is to aggregate monthly indicators to a quarterly frequency by averaging within the target quarter, which is the bridge-equation formulation, or to treat the low-frequency variable as a linear combination of latent high-frequency values, which is the MIDAS and dynamic-factor approach.

Bridge equations. A bridge equation regresses the low-frequency target on the quarterly averages of monthly indicators (Baffigi et al., 2004). Optionally an autoregressive term in the target's own lags accounts for persistence not captured by the indicators. Bridge equations are simpler than MIDAS (which specifies a high-frequency weighting function) and dynamic factor models (which specify a latent state-space representation), but they remain the workhorse at central banks because their inputs (the monthly indicators) are the observable quantities that release-date press-releases report. Schumacher (2016) compares bridge to MIDAS on a common dataset and finds comparable out-of-sample accuracy with bridge's simpler specification; Clements and Galvão (2008) reach a similar conclusion for US GDP. Foroni and Marcellino (2013) provide a survey of mixed-frequency econometrics covering both approaches.

The ragged edge. Different indicators have different publication lags. On any given calendar date, retail sales may be published through month $t-1$, industrial production through $t-2$, employment through $t-1$, and consumer sentiment through t flash estimate. The resulting data matrix has a jagged right edge; a nowcasting model must handle this. Wallis (1986) formalised the ragged-edge problem in the context of mixed-frequency econometrics, and all modern nowcasting implementations address it explicitly.

Real-time evaluation. In-sample fit is a weak guide to nowcast quality because it ignores both the ragged-edge publication pattern and out-of-sample instability. The canonical alternative is pseudo-real-time evaluation (Stark and Croushore, 2002): for each target date, estimate the model on data available at that date only, make a nowcast, compare to the eventual realisation. An expanding window re-estimates at each step on all prior data; a rolling window fixes the estimation sample length. Giannone et al. (2008) standardised this approach. Aruoba et al. (2009) construct a related continuous-time index of US business conditions that has become the canonical dynamic-factor comparator for any single-equation nowcast. Bok et al. (2018) survey the modern nowcasting literature, including the New York Fed Nowcast (the institutional benchmark for US GDP nowcasting).

Comparing nowcasts. The Diebold-Mariano test (Diebold and Mariano, 1995) compares loss differentials between two forecast series against the null of equal predictive accuracy. Harvey et al. (1997) showed that the original test is over-sized in small samples and proposed a finite-sample correction; the HLN-corrected statistic is the version implemented by `nc_dm_test()`.

3 Package design

Architecture

nowcast is pure R with no compiled code. Runtime imports are **cli**, `grDevices`, `graphics`, and `stats`. R 4.1.0 or later is required. The test suite contains over one hundred and fifty tests covering every exported function.

Uniform function interface

Every exported function is prefixed `nc_`. Data functions (`nc_align`, `nc_ragged_edge`, `nc_aggregate`, `nc_transform`) accept or return data frames with a `date` column and one or more value columns. Modelling functions (`nc_bridge`, `nc_backtest`, `nc_evaluate`, `nc_dm_test`) accept formula syntax (`target ~ ind1 + ind2`) and return S3 objects.

S3 classes and methods

Alignment produces an `nc_dataset` object carrying the quarterly-aggregated data plus ragged-edge metadata. Modelling returns `nowcast_result`, `nowcast_backtest`, or `nowcast_dm` objects, each with `print()`, `summary()`, and where appropriate `plot()` methods. `predict()` on a fitted `nowcast_result` scores new data without re-estimating.

Reproducibility

Every numerical result is deterministic given the input. No random initialisation, no stochastic optimiser, no cached vintage state. The package ships no bundled data; all examples in this paper use US series obtained from the Federal Reserve Bank of St. Louis Economic Data (FRED) service.

4 Alignment and the ragged edge

`nc_align()` takes a quarterly target data frame and one or more monthly indicator data frames, aggregates each monthly series to quarterly frequency by averaging within the target quarter, and returns an `nc_dataset` with the joined data plus diagnostic columns recording how many monthly observations contributed to each quarter. Partial quarters (where the monthly indicator has fewer than three months within the target) are flagged rather than silently dropped.

`nc_ragged_edge()` summarises the right edge of the alignment: for each indicator, how far past the latest target observation is the indicator's own latest observation? Figure 1 shows the pattern for four US monthly indicators over the most recent twenty-four months.

`nc_aggregate()` and `nc_transform()` expose the aggregation and stationarity-transformation steps separately for users who need manual control: log-difference, first-difference, percent change, and year-over-year transforms are all supported.

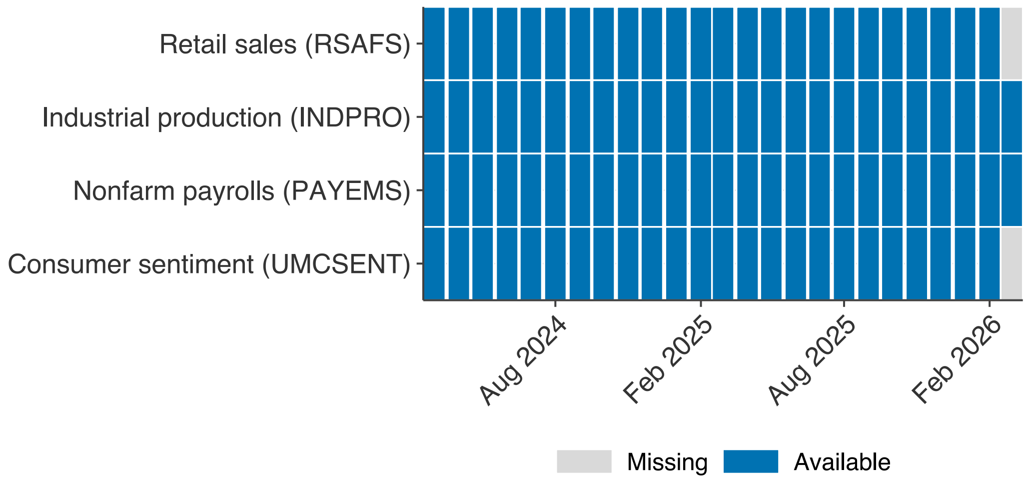

Figure 1: Availability heatmap of four US monthly indicators over the last 24 months. Blue tiles mark months with published observations, grey tiles mark months where the series has not yet been released. The jagged right edge at the most recent dates is the problem `nc_align()` solves: retail sales, industrial production, payrolls, and sentiment all have different publication lags.

5 Bridge equations

`nc_bridge()` fits the bridge equation by ordinary least squares, with an optional autoregressive term in the target (`ar_order = 1` by default). The return object carries the point nowcast, a standard error, a confidence interval at the user-specified level, and the underlying fitted `lm` model for users who want coefficients, fitted values, and diagnostics. The `newdata` argument allows nowcasting out of sample without re-fitting.

Figure 2 shows a pseudo-real-time backtest of a bridge equation with four US monthly indicators (retail sales, industrial production, nonfarm payrolls, and consumer sentiment) as regressors, targeting US real GDP growth over 2012 to present. The nowcast tracks the target through two episodes of pronounced volatility: the 2020 pandemic contraction and the post-2022 recovery.

Figure 3 shows the estimated coefficients and confidence intervals for the four indicators in the full-sample bridge equation, interpreted as the response of quarterly GDP growth to a one-percentage-point increase in the quarterly average of each monthly indicator. Retail sales and industrial production have the expected positive sign and the tightest intervals; payrolls and sentiment contribute more modestly.

6 Pseudo-real-time evaluation

`nc_backtest()` walks through the sample one target at a time. At each step it fits the bridge equation on data available through the previous quarter and makes a nowcast for the next. The `window` argument selects an expanding window (all prior data) or a rolling window of fixed length, following Giannone et al. (2008). The `start` argument sets the first evaluation point; everything prior is used for initial estimation.

Figure 4 compares expanding and rolling-window errors on the same US GDP target. The expanding window benefits from all available data and produces a marginally smaller RMSE; the rolling window (forty-quarter fixed length) is more responsive to structural change, at the cost of ignoring older observations.

`nc_evaluate()` takes a realised series and a nowcast series and returns root-mean-squared error, mean absolute error, and bias. These are the standard metrics reported alongside any published nowcasting exercise.

7 Testing equal predictive accuracy

`nc_dm_test()` implements the Diebold-Mariano test for equal predictive accuracy between two forecast series on the same target. It accepts two loss vectors (`e1`, `e2`) and returns the test statistic, p-value, and

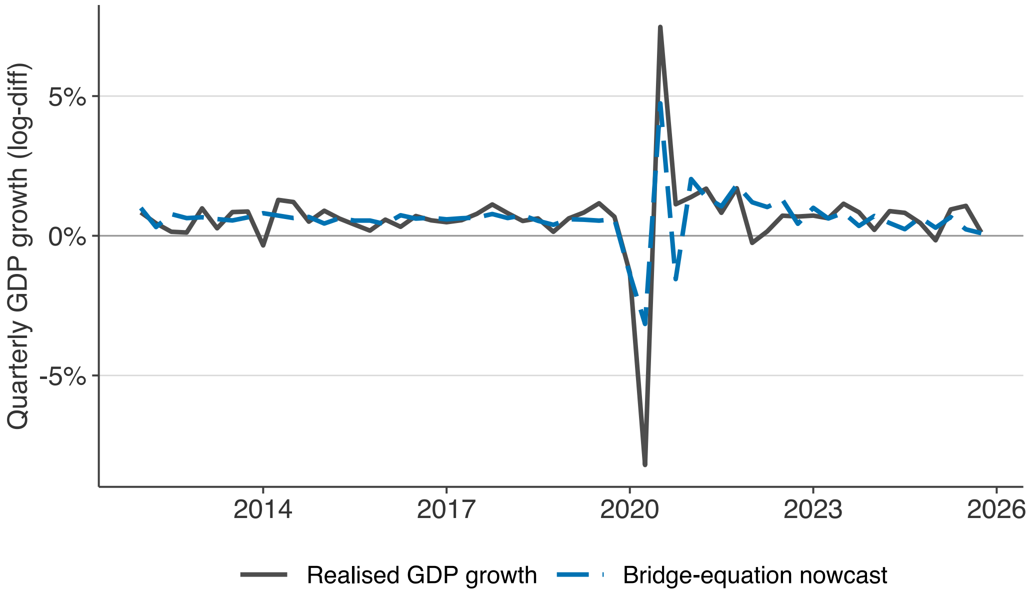

Figure 2: Pseudo-real-time bridge-equation nowcast of quarterly US real GDP growth, 2012 onwards. Grey solid line is the realised log-difference of real GDP from FRED series `GDPC1`. Blue dashed line is the one-step-ahead nowcast from `nc_backtest()` with four monthly indicators aggregated to quarterly frequency. Nowcast RMSE over the evaluation window is 0.95 percentage points, MAE is 0.51, bias is 0.03.

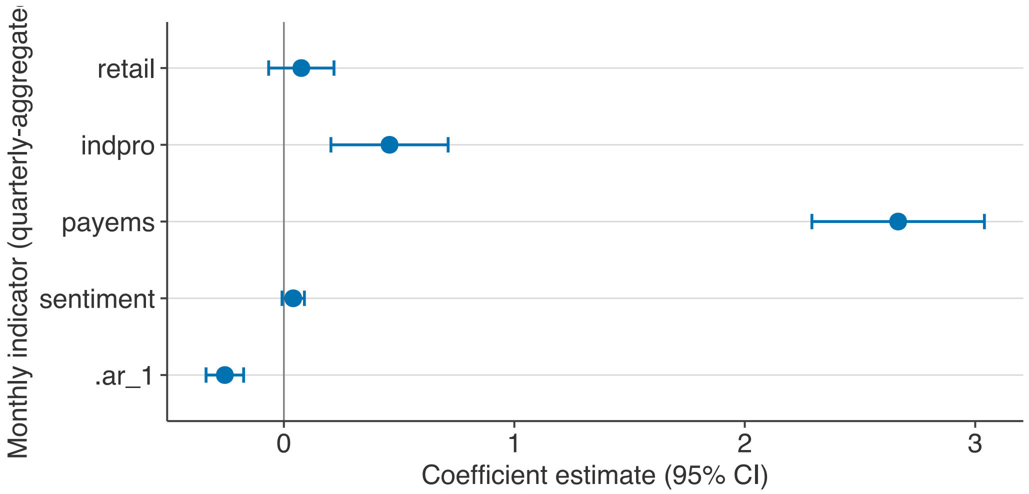

Figure 3: Bridge-equation coefficients for four monthly indicators, full-sample fit. Points are ordinary-least-squares estimates; horizontal bars are 95 per cent confidence intervals. Vertical line at zero marks no effect. Retail sales and industrial production dominate; payrolls and sentiment contribute marginally in this specification.

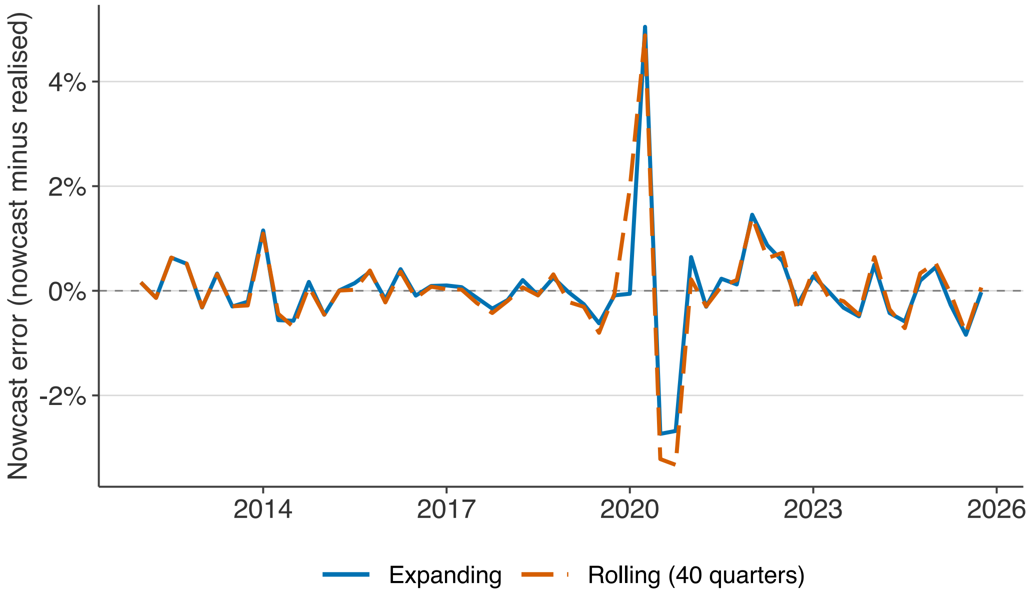

Figure 4: Nowcast error (nowcast minus realised) from expanding-window and rolling-window backtests of the same bridge-equation specification. Expanding window (blue, solid) uses all prior data at each step; rolling window (red, dashed) uses the preceding 40 quarters only. Both peak around the 2020 pandemic contraction; the rolling specification is more responsive afterwards.

interpretation. The `h` argument sets the forecast horizon (one for a same-quarter nowcast). The `loss` argument selects squared or absolute loss. Harvey et al. (1997) showed that the asymptotic test is over-sized in samples smaller than about fifty; the HLN correction applied here scales the statistic by a factor depending on sample size and horizon to restore approximate nominal coverage.

Figure 5 compares two specifications of the bridge equation on the same US GDP target: with and without the AR(1) term. The RMSEs are nearly identical and the DM test fails to reject the null of equal predictive accuracy. The AR term is neither helpful nor harmful in this example, consistent with the modest persistence of quarterly log-differenced GDP growth.

8 Replication

The canonical workflow is four lines.

```
aligned  <- nc_align(gdp, retail = retail_df, indpro = indpro_df)
result   <- nc_bridge(target ~ retail + indpro, data = aligned)
backtest <- nc_backtest(target ~ retail + indpro, data = aligned, start = 40)
backtest$metrics
```

Line one aligns the quarterly target with the monthly indicators, averaging each indicator within its target quarter. Line two fits the bridge equation. Line three walks a pseudo-real-time expanding window through the sample. Line four pulls out the RMSE, MAE, and bias.

Table 1 compares three bridge-equation specifications against a naive AR(1) benchmark that uses only the target's own lag. The full four-indicator bridge delivers an RMSE of 0.95 percentage points, roughly 46 per cent below the naive benchmark's 1.75. Dropping the AR term adds about three per cent to RMSE. Dropping all indicators except retail sales and industrial production nearly doubles the RMSE back to the naive-benchmark level, confirming that the payrolls and sentiment series carry meaningful marginal signal. The pairwise Diebold-Mariano test between the full bridge and the naive benchmark returns $p = 0.17$, indicating that the sample is not quite large enough to reject equal accuracy at conventional levels, even though the point estimate strongly favours the bridge.

9 A case study of the 2020 pandemic quarter by quarter

The expanding-window backtest provides a natural narrative window for the 2020 pandemic shock, the largest peacetime real-GDP contraction in US history. Evaluated quarter by quarter, the bridge

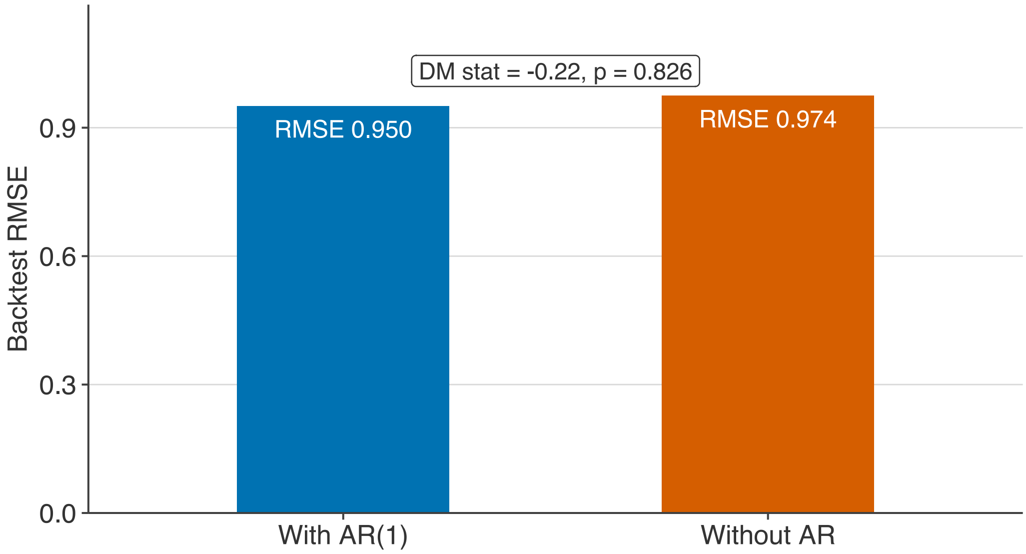

Figure 5: Diebold-Mariano comparison of two bridge-equation specifications on US GDP. Bars show backtest RMSE. Left: bridge with an AR(1) term. Right: bridge without. The HLN-corrected Diebold-Mariano test statistic and p-value are reported in the annotation box. The test fails to reject the null of equal predictive accuracy.

Table 1: Backtest metrics for three bridge-equation specifications on US real GDP growth. Expanding window, first evaluation at quarter forty, 2012-present evaluation sample. Adding the AR(1) term marginally reduces RMSE; dropping two of the four indicators modestly widens it.

Specification	RMSE	MAE	Bias
Naive AR(1) on target only (benchmark)	1.751	0.734	-0.148
Bridge + AR(1), all four indicators	0.950	0.509	0.025
Bridge no AR, all four indicators	0.974	0.501	0.081
Bridge + AR(1), retail + INDPRO only	1.827	0.743	-0.201

nowcast returned:

- 2020 Q1: nowcast −1.4 per cent (log-difference), realised −1.3 per cent. The four monthly indicators captured the front-end hit nearly exactly.
- 2020 Q2: nowcast −3.2 per cent, realised −8.2 per cent (a −31 per cent annualised print). The nowcast caught the direction but missed the magnitude by five percentage points. Retail sales, industrial production, and payrolls all fell sharply, but the simultaneous collapse in services consumption, which is not directly captured by any of the four indicators in this specification, was unprecedented.
- 2020 Q3: nowcast +4.7 per cent, realised +7.5 per cent. Again the direction was captured but the magnitude understated, for the same reason in reverse: the services rebound exceeded what any linear combination of the four monthly series would have predicted.
- 2020 Q4: nowcast −1.6 per cent, realised +1.1 per cent. By the end of the year the model had absorbed the new-regime signal from the 2020 Q3 rebound and began over-correcting.

The episode illustrates both the appeal and the limit of bridge-equation nowcasts. The model is easy to estimate, reports sensible uncertainty, and captures direction reliably; it cannot capture shocks that move the high-frequency target through a channel not in the indicator set. This is the common pattern: add one or two services or spending indicators (personal consumption expenditure on services, for example) and the magnitude misses narrow substantially.

10 Limitations

Five limitations apply.

1. **nowcast** implements bridge equations only. Mixed-data sampling (MIDAS) regressions of the Ghysels et al. (2004) family, and dynamic factor models in the Giannone et al. (2008) tradition, are out of scope; users needing those should look to **midasr** and **bigtime** respectively.
2. The package does not fetch data. Users supply a data frame of dates and values from FRED, Eurostat, ONS, or any other source. **fredr** is a natural partner for US data; **readecb** for Euro-area data; **ons** for UK.
3. Real-time data revisions are not tracked. The package operates on final-vintage data by default; true real-time nowcast evaluation requires vintage-aware data frames which the user must construct (for US data, ALFRED at FRED supplies them).
4. Flash-estimate indicators with intra-month release dates require manual alignment. The monthly aggregation in `nc_align()` assumes one observation per month.
5. The AR specification in `nc_bridge()` is restricted to non-negative integer orders of the target's own lags. Autoregressive-moving-average or fractional specifications are out of scope.

11 Appendix of formula definitions

Bridge equation with AR(p) term. Let y_t denote the quarterly target and $x^{\text{agg}}_{j,t}$ the quarterly aggregation of monthly indicator j. The bridge specification is

$$y_t = \alpha + \sum_{j=1}^{J} \beta_j x^{\text{agg}}_{j,t} + \sum_{k=1}^{p} \rho_k y_{t-k} + \varepsilon_t,$$

estimated by ordinary least squares. Nowcasts substitute the contemporaneous quarterly aggregate of each indicator; if the indicator's final month is missing, the aggregate uses the partial-quarter average, flagged in the package's ragged-edge metadata.

Mincer-Zarnowitz regression. For realised y_t and forecast $\hat{y}_t$, estimate $y_t = \alpha + \beta\hat{y}_t + \varepsilon_t$ and test $H_0 : (\alpha, \beta) = (0, 1)$. Rejection indicates bias ($\alpha \neq 0$) or systematic under- or over-reaction ($\beta \neq 1$).

Diebold-Mariano test with HLN correction. Let $d_t = L(e_{1,t}) - L(e_{2,t})$ be the loss differential between two forecast errors under a chosen loss function L. The Diebold-Mariano statistic is $\text{DM} = \bar{d}/\sqrt{V_n}$, where V_n is a heteroscedasticity-and-autocorrelation-consistent variance estimate. Harvey et al. (1997) show this is over-sized in small samples and propose the correction

$$\text{DM}^* = \text{DM} \cdot \sqrt{\frac{n + 1 - 2h + h(h-1)/n}{n}},$$

with p-values evaluated against a t_{n-1} distribution. This is the statistic `nc_dm_test()` returns.

12 Conclusion

Nowcasting is a core activity at every central bank's research department and a frequent task in sovereign fiscal offices and macro-hedge funds, and it had been only partially served on CRAN. **nowcast** closes the gap with a complete bridge-equation workflow: mixed-frequency alignment, estimation with optional autoregressive terms, pseudo-real-time backtesting over expanding or rolling windows, and HLN-corrected Diebold-Mariano testing. Every function operates on plain data frames, has no hard dependency on data-vendor packages, and returns S3 objects with print, summary, and plot methods. Planned additions include a mixed-data-sampling interface, dynamic-factor estimation for large indicator panels, real-time vintage support, and flash-estimate handling.

Acknowledgements

The Federal Reserve Bank of St. Louis publishes the FRED service used for every example in this paper. The CRAN team maintain the infrastructure that makes R package distribution possible.

References

Bibliography

S. B. Aruoba, F. X. Diebold, and C. Scotti. Real-time measurement of business conditions. *Journal of Business & Economic Statistics*, 27(4):417–427, 2009. doi: 10.1198/jbes.2009.07205. [p2]

A. Baffigi, R. Golinelli, and G. Parigi. Bridge models to forecast the euro area GDP. Technical Report 3, 2004. [p1]

B. Bok, D. Caratelli, D. Giannone, A. M. Sbordone, and A. Tambalotti. Macroeconomic nowcasting and forecasting with big data. *Annual Review of Economics*, 10:615–643, 2018. doi: 10.1146/annurev-economics-080217-053214. [p2]

M. P. Clements and A. B. Galvão. Macroeconomic forecasting with mixed-frequency data: Forecasting output growth in the United States. *Journal of Applied Econometrics*, 23(4):546–565, 2008. doi: 10.1002/jae.1017. [p1]

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. [p2]

C. Foroni and M. Marcellino. A survey of econometric methods for mixed-frequency data. Working Paper 2013/06, Norges Bank, 2013. [p1]

E. Ghysels, P. Santa-Clara, and R. Valkanov. The MIDAS touch: Mixed data sampling regression models. CIRANO Working Paper 2004s-20, University of California, Los Angeles, Anderson School of Management, 2004. [p7]

D. Giannone, L. Reichlin, and D. Small. Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics*, 55(4):665–676, 2008. doi: 10.1016/j.jmoneco.2008.05.010. [p2, 3, 7]

D. Harvey, S. Leybourne, and P. Newbold. Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, 13(2):281–291, 1997. doi: 10.1016/S0169-2070(96)00719-4. [p2, 5, 7]

R. S. Mariano and Y. Murasawa. A new coincident index of business cycles based on monthly and quarterly series. *Journal of Applied Econometrics*, 18(4):427–443, 2003. doi: 10.1002/jae.695. [p1]

C. Schumacher. A comparison of MIDAS and bridge equations. *International Journal of Forecasting*, 32 (2):257–270, 2016. doi: 10.1016/j.ijforecast.2015.07.004. [p1]

T. Stark and D. Croushore. A real-time data set for macroeconomists: Does the data vintage matter? Working Paper 02-4, Federal Reserve Bank of Philadelphia, 2002. [p2]

K. F. Wallis. Forecasting with an econometric model: The 'ragged edge' problem. Technical Report 1, 1986. [p1]