Package 'forecastdom'

November 1, 2026

Title (Un)Conditional Forecast Dominance

Version 0.1.1

Description A unified toolkit for out-of-sample forecast dominance testing. Covers unconditional and conditional equal and superior predictive ability, encompassing, and nested-model comparison. Implements the Diebold-Mariano test with the Harvey, Leybourne, and Newbold (1997) <doi:10.1016/S0169-2070(96)00719-4> small-sample correction; the Clark-West mean squared forecast error (MSFE) adjusted statistic (Clark and West, 2007) <doi:10.1016/j.jeconom.2006.05.023>; the ENC-NEW encompassing test of Clark and McCracken (2001) <doi:10.1016/S0304-4076(01)00071-9>; the Giacomini-White conditional equal predictive ability test (Giacomini and White, 2006) <doi:10.1111/j.1468-0262.2006.00718.x>; Hansen's superior predictive ability (SPA) test (Hansen, 2005) <doi:10.1198/073500105000000063>; the conditional superior predictive ability test of Li, Liao, and Quaedvlieg (2022) <doi:10.1093/restud/rdab039>; and the uniform and average multi-horizon SPA tests of Quaedvlieg (2021) <doi:10.1080/07350015.2019.1620074>.

License MIT + file LICENSE

Encoding UTF-8

Language en-GB

RoxygenNote 7.3.2

LinkingTo Rcpp

Imports MASS, Rcpp, stats, ggplot2

Suggests knitr, R.matlab, rmarkdown, testthat (>= 3.0.0)

VignetteBuilder knitr

URL https://gabbocg.github.io/forecastdom/, https://github.com/gabbocg/forecastdom

BugReports https://github.com/gabbocg/forecastdom/issues

Config/testthat/edition 3

Depends R (>= 3.5)

LazyData true

NeedsCompilation yes

Author Gabriel Cabrera [aut, cre] (ORCID:
<https://orcid.org/0000-0001-8803-2338>)

Maintainer Gabriel Cabrera <gabriel.cabrera.guz@gmail.com>

Repository CRAN

Date/Publication 2026-11-01 17:20:02 UTC

Contents

aspa_mh_test . . . 3
cm2001 . . . 4
csms . . . 5
cspa_test . . . 6
cspa_test_plot . . . 8
cw_test . . . 9
dm_test . . . 10
do_sim . . . 11
enc_new . . . 12
gw2006 . . . 13
gw_test . . . 14
hl2005 . . . 15
ivx_wald . . . 16
llq2022 . . . 17
llq2022_jnj . . . 18
llq2022_uv_cspa . . . 19
mse_f_test . . . 20
nrtz2014 . . . 21
print.cspa_test . . . 22
qll_hat . . . 23
quaedvlieg2021 . . . 24
rossi2006 . . . 25
rrz2016 . . . 26
rz2013 . . . 27
spa_test . . . 28
summary.cspa_test . . . 29
uspa_mh_test . . . 30
wg2008 . . . 31

Index **33**

`aspa_mh_test` *Average Multi-Horizon Superior Predictive Ability Test*

Description

Implements the average multi-horizon SPA test of Quaedvlieg (2021), which tests whether a weighted average of horizon-wise expected loss differentials favours the benchmark, allowing inferior performance at some horizons to be compensated by superior performance at others. Bootstrap follows Algorithm 1 (moving-block bootstrap, block length `L`).

Usage

```
aspa_mh_test(loss_diff, weights, L, B = 999L, level = 0.05)
```

Arguments

`loss_diff`	A `T` x `H` matrix of loss differentials, defined as `L_bench - L_comp` at each horizon. Positive values indicate the benchmark has higher loss (i.e. is worse) than the competitor.
`weights`	Numeric vector of length `H` of horizon weights. They need not sum to 1; uniform weights `rep(1/H, H)` reproduce the unweighted average SPA of the original paper.
`L`	Integer block length for the moving-block bootstrap.
`B`	Integer number of bootstrap replications. Default `999`.
`level`	Significance level. Default `0.05`.

Details

The aSPA statistic is

$$t^{aSPA} = \sqrt{T}\,\bar{d}^{w} / \sqrt{\hat{\omega}^{QS}(\bar{d}^{w})},$$

where $d_t^w = \sum_h w_h d_{h,t}$ and $\hat{\omega}^{QS}$ is the Quadratic-Spectral HAC long-run variance (Andrews, 1991). Critical values are obtained via a moving-block bootstrap that recentres the weighted series under the null. The p-value is upper-tail (`mean(t < t_b)`); rejection requires the weighted-average differential to be large and positive (the benchmark loses on the weighted average).

This is a port of the Matlab reference implementation accompanying Quaedvlieg (2021); the helpers `.mbb_indices`, `.mbb_variance`, and `.qs_lrvar` translate `Get_MBB_ID.m`, `MBB_Variance.m`, and `QS.m` respectively.

Value

A list with class `"aspa_mh_test"` containing `statistic`, `pvalue`, `reject`, `level`, `L`, `B`, `T`, `H`, `weights`, `d_bar`.

References

Quaedvlieg, R. (2021). Multi-Horizon Forecast Comparison. *Journal of Business & Economic Statistics*, 39(1), 40-53.

See Also

uspa_mh_test, spa_test

Examples

```
set.seed(2)
ld <- matrix(rnorm(200 * 4, mean = 0.2), 200, 4)
aspa_mh_test(ld, weights = rep(0.25, 4), L = 3, B = 199)
```

---

cm2001 *US Unemployment and Inflation (Clark & McCracken, 2001 setup)*

Description

Monthly US civilian unemployment rate and annualised monthly log inflation rate, assembled from FRED series UNRATE and CPIAUCSL. The series support the Phillips-curve forecasting illustration in Clark and McCracken (2001, Section 5): does adding lagged inflation to a univariate AR for unemployment improve out-of-sample forecasts?

Usage

```
cm2001
```

Format

A data frame with 937 rows and 3 columns:

date First day of the month (Date).

unrate Civilian unemployment rate (percent), FRED UNRATE.

infl Annualised monthly log inflation, computed as $1200 \cdot \log(\text{CPI}_t / \text{CPI}_{t-1})$ from FRED CPIAUCSL.

Source

FRED, Federal Reserve Bank of St. Louis, https://fred.stlouisfed.org/series/UNRATE and https://fred.stlouisfed.org/series/CPIAUCSL.

References

Clark, T. E. and McCracken, M. W. (2001). Tests of equal forecast accuracy and encompassing for nested models. *Journal of Econometrics*, 105(1), 85-110.

csms *Confidence Set for the Most Superior (CSMS) Forecasting Method*

Description

Constructs the confidence set for the most superior forecasting method by inverting the CSPA test, as described in Section 2.3 of Li, Liao, and Quaedvlieg (2022). The set contains all methods `j` for which the CSPA null hypothesis (with `j` as benchmark) is not rejected.

Usage

```
csms(
  losses,
  X,
  level,
  trim = 0,
  prewhiten = -1L,
  preselect = TRUE,
  R = 10000L,
  method_names = NULL
)
```

Arguments

`losses`	An `n x (J+1)` matrix where each column contains the loss series for a forecasting method. All methods are treated symmetrically (no pre-specified benchmark).
`X`	An `n x 1` numeric vector of the conditioning variable.
`level`	Significance level (e.g., 0.05).
`trim`	Trimming parameter (standard deviations). Default `0`.
`prewhiten`	Pre-whitening order. Default `-1` (AIC).
`preselect`	Logical; adaptive inequality selection. Default `TRUE`.
`R`	Integer; Gaussian process replications. Default `10000`.
`method_names`	Optional character vector of method names. If `NULL`, uses `"M1"`, `"M2"`, ...

Details

For each method `j` in `0, ..., J`, the CSPA test is applied with `j` as the benchmark. The confidence set is:

$$\widehat{\mathcal{M}}_{n,1-\alpha} = \{0 \leq j \leq J : \text{CSPA test with } j \text{ as benchmark does not reject}\}$$

Value

A list with class "csms" containing:

in_set	Logical vector; which methods are in the confidence set.
set_members	Names of methods in the confidence set.
theta	Theta values for each method as benchmark.
pvalues	P-values for each method as benchmark.
method_names	Names of all methods.
level	Significance level.
n_methods	Total number of methods.

References

Li, J., Liao, Z., and Quaedvlieg, R. (2022). Conditional Superior Predictive Ability. *Review of Economic Studies*, 89(2), 843-875.

Examples

```
set.seed(42)
n <- 300
X <- arima.sim(list(ar = 0.5), n = n, sd = sqrt(0.75))
losses <- matrix(rnorm(n * 4), n, 4)
csms(losses, X, level = 0.05, trim = 2, R = 500L)
```

---

cspa_test *Conditional Superior Predictive Ability (CSPA) Test*

Description

Tests the null hypothesis that a benchmark forecasting method has conditional superior predictive ability over competing alternatives, uniformly across all conditioning states. Based on Li, Liao, and Quaedvlieg (2022).

Usage

```
cspa_test(Y, X, level, trim = 0, prewhiten = -1L, preselect = TRUE, R = 10000L)
```

Arguments

Y	An n x J matrix of loss differentials with respect to the benchmark. Positive values indicate the benchmark outperforms competitor j in period t.
X	An n x 1 numeric vector of the conditioning variable.
level	Significance level (e.g., 0.05).

`trim`	Trim observations where the conditioning variable exceeds this many standard deviations from the mean. Use `0` (default) for no trimming.
`prewhiten`	Order of pre-whitening VAR for HAC estimation. Use `0` for standard Newey-West, or `-1` (default) for AIC-based lag selection.
`preselect`	Logical; perform adaptive inequality selection? Default `TRUE`.
`R`	Integer; number of bootstrap replications for critical value computation. Default `10000`.

Value

A list with class `"cspa_test"` containing:

`theta`	Infimum of the upper confidence bound. Negative values lead to rejection of the null.
`pvalue`	P-value for the test.
`reject`	Logical; whether the null is rejected at the given level.
`level`	Significance level used.
`h_hat`	Estimated conditional mean functions (`n` x `J` matrix).
`sigma_jx`	Estimated standard deviations (`J` x `n` matrix).
`kp`	Critical value from adaptive inequality selection.
`Vhat`	Logical matrix indicating selected (j, x) pairs.
`X`	Conditioning variable (after trimming).
`Y`	Loss differentials (after trimming).
`K`	Number of series terms used.
`prewhiten_order`	Pre-whitening lag order actually used.

References

Li, J., Liao, Z., and Quaedvlieg, R. (2022). Conditional Superior Predictive Ability. *Review of Economic Studies*, 89(2), 843-875.

Examples

```
sim <- do_sim(J = 3, n = 250, a = 1, c = 0, rho_u = 0.4)
result <- cspa_test(sim$Y, sim$X, level = 0.05, trim = 2, R = 500L)
print(result)
```

cspa_test_plot *Plot CSPA Test Results*

Description

Visualizes the estimated conditional mean functions $\hat{h}_j(x)$, their lower envelope, and the confidence bound from the CSPA test.

Usage

```
cspa_test_plot(
  object = NULL,
  Y = NULL,
  X = NULL,
  level = 0.05,
  trim = 0,
  prewhiten = -1L,
  preselect = TRUE,
  ylab = "Loss Difference",
  xlab = "Conditioning Variable"
)
```

Arguments

object	A cspa_test object returned by cspa_test, or omitted if Y and X are provided directly.
Y	An n x J matrix of loss differentials. Ignored if object is provided.
X	An n x 1 numeric vector of the conditioning variable. Ignored if object is provided.
level	Significance level (e.g., 0.05). Ignored if object is provided.
trim	Trimming parameter. Default 0.
prewhiten	Pre-whitening order. Default -1.
preselect	Logical; adaptive inequality selection. Default TRUE.
ylab	Label for the y-axis. Default "Loss Difference".
xlab	Label for the x-axis. Default "Conditioning Variable".

Value

A ggplot object.

`cw_test` *Clark-West Test for Predictive Ability of Nested Models*

Description

Tests whether an alternative (unrestricted) model has superior out-of-sample predictive ability relative to a benchmark (restricted) model, using the MSFE-adjusted statistic of Clark and West (2007). Also computes the out-of-sample R^2_{OS} statistic.

Usage

```
cw_test(e1, e2, f1, f2)
```

Arguments

`e1`	Numeric vector of forecast errors from the benchmark (restricted/null) model.
`e2`	Numeric vector of forecast errors from the alternative (unrestricted) model.
`f1`	Numeric vector of forecasts from the benchmark model.
`f2`	Numeric vector of forecasts from the alternative model.

Details

The MSFE-adjusted series is defined as:

$$\hat{f}_t = e_{1,t}^2 - \left(e_{2,t}^2 - (f_{1,t} - f_{2,t})^2\right)$$

The test regresses $\hat{f}_t$ on a constant and uses the resulting t-statistic, compared to a standard normal distribution (one-sided).

Value

A list with class `"cw_test"` containing:

`statistic`	The Clark-West t-statistic.
`pvalue`	One-sided p-value (H1: alternative is better).
`r2os`	Out-of-sample R^2_{OS} in percent.
`n`	Number of observations.

References

Clark, T.E. and West, K.D. (2007). Approximately Normal Tests for Equal Predictive Accuracy in Nested Models. *Journal of Econometrics*, 138(1), 291-311.

Examples

```
set.seed(42)
n <- 200
actual <- rnorm(n)
f1 <- actual + rnorm(n, sd = 0.5)
f2 <- actual + rnorm(n, sd = 0.4)
e1 <- actual - f1
e2 <- actual - f2
cw_test(e1, e2, f1, f2)
```

---

dm_test *Diebold-Mariano Test for Equal Predictive Ability*

Description

Tests the null hypothesis that two forecasting methods have equal predictive ability (UEPA), with an optional small-sample correction by Harvey, Leybourne, and Newbold (1997).

Usage

```
dm_test(
  e1,
  e2,
  h = 1L,
  loss = c("SE", "AE"),
  alternative = c("two.sided", "less", "greater"),
  correction = TRUE
)
```

Arguments

e1	Numeric vector of forecast errors from model 1 (benchmark).
e2	Numeric vector of forecast errors from model 2 (competitor).
h	Integer; forecast horizon. Default 1.
loss	Character; loss function to use. "SE" for squared error (default), "AE" for absolute error.
alternative	Character; alternative hypothesis. "two.sided" (default), "less" (model 2 is better), or "greater" (model 1 is better).
correction	Logical; apply the Harvey, Leybourne, and Newbold (1997) finite-sample correction? Default TRUE.

Value

A list with class "dm_test" containing:

statistic	The (possibly corrected) DM test statistic.
pvalue	P-value.
alternative	The alternative hypothesis used.
correction	Whether HLN correction was applied.
h	Forecast horizon.
n	Number of observations.
loss	Loss function used.

References

Diebold, F.X. and Mariano, R.S. (1995). Comparing Predictive Accuracy. *Journal of Business & Economic Statistics*, 13(3), 253-263.

Harvey, D., Leybourne, S., and Newbold, P. (1997). Testing the Equality of Prediction Mean Squared Errors. *International Journal of Forecasting*, 13(2), 281-291.

Examples

```
set.seed(42)
e1 <- rnorm(100)
e2 <- rnorm(100, mean = 0.1)
dm_test(e1, e2)
```

---

do_sim *Simulate Data from the CSPA Paper DGP*

Description

Generates data from the data generating process in Section 3.1 of Li, Liao, and Quaedvlieg (2022). The conditioning variable X_t follows a Gaussian AR(1) with coefficient 0.5, and the loss differentials are $Y_{j,t} = 1 - a\exp(-(X_t - c)^2) + u_{j,t}$, where $u_{j,t}$ follows an AR(1) with coefficient rho_u.

Usage

```
do_sim(J, n, a, c, rho_u)
```

Arguments

J	Integer; number of competing forecast methods.
n	Integer; sample size.
a	Numeric; controls the shape of the conditional mean function. a = 1 gives the null hypothesis ($h_j(c) = 0$).
c	Numeric; location parameter for the minimum of $h_j(x)$.
rho_u	Numeric in $[0, 1)$; AR(1) coefficient of the error process. Higher values induce more serial dependence.

Value

A list with components:

Y	An n x J matrix of loss differentials.
X	A numeric vector of length n.

Examples

```
sim <- do_sim(J = 3, n = 250, a = 1, c = 0, rho_u = 0.4)
str(sim)
```

---

enc_new *Clark-McCracken ENC-NEW Encompassing Test*

Description

Tests whether the benchmark model encompasses the alternative model, using the ENC-NEW statistic of Clark and McCracken (2001). Under the null, the benchmark forecast encompasses the alternative (no additional information in the alternative).

Usage

```
enc_new(e1, e2)
```

Arguments

e1	Numeric vector of forecast errors from the benchmark model.
e2	Numeric vector of forecast errors from the alternative model.

Details

The ENC-NEW statistic is:

$$\text{ENC-NEW} = n \cdot \frac{\bar{c}}{\text{MSFE}_2}$$

where $\bar{c} = n^{-1} \sum (e_{1,t}^2 - e_{2,t} e_{1,t})$ and $\text{MSFE}_2 = n^{-1} \sum e_{2,t}^2$.

Critical values are non-standard and depend on the estimation setup. See Clark and McCracken (2001, Table 2) for asymptotic critical values.

Value

A list with class "enc_new" containing:

statistic	The ENC-NEW test statistic.
n	Number of observations.

References

Clark, T.E. and McCracken, M.W. (2001). Tests of Equal Forecast Accuracy and Encompassing for Nested Models. *Journal of Econometrics*, 105(1), 85-110.

Examples

```
set.seed(42)
e1 <- rnorm(100)
e2 <- rnorm(100, sd = 0.9)
enc_new(e1, e2)
```

---

gw2006 *SPF Mean CPI Inflation Forecasts (Giacomini & White, 2006 setup)*

Description

Quarterly mean Survey of Professional Forecasters (SPF) forecasts of annualised CPI inflation at horizons $h = 0, 1, 2, 3, 4$ quarters ahead, aligned with realised CPI inflation from FRED. The panel covers 1981Q3 onward, the start of the SPF's CPI survey.

Usage

```
gw2006
```

Format

A data frame with 180 rows and 10 columns:

date First day of the third month of the realisation quarter (Date).

year, quarter Realisation quarter.

infl Realised annualised QoQ CPI inflation (percent), computed from quarterly averages of CPIAUCSL.

infl_lag One-quarter lag of infl — the naive random-walk forecast benchmark.

spf_h0, spf_h1, spf_h2, spf_h3, spf_h4 SPF mean forecasts of infl, made $h = 0, \ldots, 4$ quarters in advance.

Source

SPF mean CPI inflation forecasts from the Federal Reserve Bank of Philadelphia, https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/survey-of-professional-forecasters. Realised CPI from FRED series CPIAUCSL, https://fred.stlouisfed.org/series/CPIAUCSL.

References

Giacomini, R. and White, H. (2006). Tests of conditional predictive ability. *Econometrica*, 74(6), 1545-1578.

`gw_test` *Conditional Equal Predictive Ability (CEPA) Test*

Description

Tests the null hypothesis of conditional equal predictive ability using the Giacomini and White (2006) approach. The test checks whether loss differentials are conditionally mean-zero given an information set, using a finite set of instruments.

Usage

```
gw_test(e1, e2, instruments = NULL, loss = c("SE", "AE"))
```

Arguments

`e1`	Numeric vector of forecast errors from model 1 (benchmark).
`e2`	Numeric vector of forecast errors from model 2 (competitor).
`instruments`	An `n` x `q` matrix of instruments (conditioning variables) observed at the time of the forecast. If `NULL` (default), uses a constant and the lagged loss differential.
`loss`	Character; loss function. `"SE"` for squared error (default), `"AE"` for absolute error.

Details

The test is based on the unconditional moment conditions implied by the CEPA null hypothesis:

$$E[d_t \cdot W_t] = 0$$

where d_t is the loss differential and W_t is a vector of instruments measurable with respect to the conditioning information set.

The test statistic follows a chi-squared distribution with q degrees of freedom, where q is the number of instruments.

Value

A list with class `"gw_test"` containing:

`statistic`	The Wald test statistic.
`pvalue`	P-value from the chi-squared distribution.
`df`	Degrees of freedom (number of instruments).
`n`	Number of observations used.
`loss`	Loss function used.

References

Giacomini, R. and White, H. (2006). Tests of Conditional Predictive Ability. *Econometrica*, 74(6), 1545-1578.

Examples

```
set.seed(42)
e1 <- rnorm(200)
e2 <- rnorm(200, sd = 0.9)
gw_test(e1, e2)
```

`hl2005` *IBM Volatility Forecasts and Realized-Variance Proxies (Hansen & Lunde, 2005)*

Description

Daily IBM data from the replication archive of Hansen and Lunde (2005). Contains 254 trading days (1999-06-01 to 2000-05-31), 8 realized-variance proxies of varying quality, and 330 one-step-ahead conditional variance forecasts produced by GARCH-family models. The benchmark GARCH(1,1) with constant mean and Gaussian errors is identified by the index `garch11_idx`.

Usage

```
hl2005
```

Format

A list with the following components:

date Trading-day `Date` vector (length 254).

rv Numeric vector (length 254); the 5-minute linear-interpolation realized-variance proxy. The paper's headline series.

rv_proxies 254×8 matrix of alternative RV proxies: `sq_ccr` (squared close-to-close returns), `spline_50_3min`, `spline_250_2min`, `fourier_M85`, `linear_5min`, `prevtick_5min`, `linear_1min`, `prevtick_1min`.

forecasts 254×330 matrix of one-step-ahead conditional variance forecasts. Columns index a base GARCH specification (55 specs: LGARCH, IGARCH, TS-GARCH, A-GARCH, NA-GARCH, V-GARCH, THR-GARCH, GJR-GARCH, LOG-GARCH, EGARCH, NGARCH, A-PARCH, GQ-ARCH, H-GARCH, AUG-GARCH) crossed with mean-equation and error-distribution combinations (zero/const/ GARCH-in-mean $\times$ Gaussian/t-distributed).

garch11_idx Integer (= 57); the column of `forecasts` corresponding to GARCH(1,1) with constant mean and Gaussian errors, used as the benchmark in the paper.

Source

Journal of Applied Econometrics data archive for Hansen and Lunde (2005), volume 20, issue 7.

References

Hansen, P. R. and Lunde, A. (2005). A forecast comparison of volatility models: does anything beat a GARCH(1,1)? *Journal of Applied Econometrics*, 20(7), 873-889.

`ivx_wald` *IVX-Wald Test for Predictive Regressions*

Description

Computes the IVX-Wald statistic of Kostakis, Magdalinos, and Stamatogiannis (2015) for testing predictability in a regression of returns on persistent predictors. The IVX approach is robust to the degree of persistence of the regressors (stationary, local-to-unity, or unit root).

Usage

```
ivx_wald(y, X, K = 1L, M_n = 0L, beta = 0.95)
```

Arguments

`y`	Numeric vector of length `T`; the dependent variable (e.g., returns).
`X`	A `T` x `r` matrix of predictor observations.
`K`	Integer; forecast horizon. Default `1`.
`M_n`	Integer; bandwidth parameter for the long-run covariance estimator. Default `0` (no correction). Use `floor(T^(1/3))` as a rule of thumb.
`beta`	Numeric in $(0, 1)$; controls the rate of the IVX instrument. Values close to 1 yield best performance. Default `0.95`.

Details

The IVX-Wald test constructs an endogenous instrument $\tilde{Z}_t$ by filtering the predictor increments through a mildly integrated process with autoregressive root $R_{nz} = 1 - 1/n^{\beta}$. The resulting Wald statistic is asymptotically chi-squared with `r` degrees of freedom, regardless of the persistence of the predictors.

Value

A list with class `"ivx_wald"` containing:

`statistic`	The IVX-Wald test statistic.
`pvalue`	P-value from the chi-squared distribution.
`coefficients`	IVX coefficient estimates.
`K`	Forecast horizon.
`n`	Number of observations.
`r`	Number of predictors.

References

Kostakis, A., Magdalinos, T., and Stamatogiannis, M.P. (2015). Robust Econometric Inference for Stock Return Predictability. *Review of Financial Studies*, 28(5), 1506-1553.

Examples

```
set.seed(42)
n <- 200
x <- cumsum(rnorm(n))
y <- 0.02 * x + rnorm(n)
ivx_wald(y, as.matrix(x))
```

`llq2022` *S&P 500 Realized-Variance Forecasts (Li, Liao & Quaedvlieg, 2022)*

Description

Daily 5-minute realized variance of the S&P 500 with one-step-ahead out-of-sample forecasts from six competing models, plus lagged VIX as a conditioning variable. The forecasts and realized variances are from the replication package of Li, Liao and Quaedvlieg (2022) on Zenodo; the underlying RV series is from Bollerslev, Patton and Quaedvlieg (2016).

Usage

```
llq2022
```

Format

A data frame with 3,071 rows and 9 variables. Sample period is 2001-05-11 to 2013-08-30 (daily, trading days).

date Trading-day `Date`.

rv Realized variance of the S&P 500 (5-minute returns), in percent-squared units.

AR1 AR(1) forecast of `rv`.

AR22 AR(22) forecast of `rv`.

AR22_Lasso AR(22) forecast with adaptive lasso (Matlab).

HAR HAR forecast of Corsi (2009).

HARQ HARQ forecast of Bollerslev, Patton and Quaedvlieg (2016).

ARFIMA Fractionally integrated AR forecast.

vix_lag VIX close from the previous trading day, used as conditioning variable.

Source

Replication package of Li, Liao and Quaedvlieg (2022), https://zenodo.org/record/4884813. Realized variance series adapted from Bollerslev, Patton and Quaedvlieg (2016), retrieved from Quaedvlieg's website. VIX from CBOE.

References

Bollerslev, T., Patton, A. J. and Quaedvlieg, R. (2016). Exploiting the errors: a simple approach for improved volatility forecasting. *Journal of Econometrics*, 192(1), 1-18.

Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2), 174-196.

Li, J., Liao, Z. and Quaedvlieg, R. (2022). Conditional Superior Predictive Ability. *Review of Economic Studies*, 89(2), 843-875.

Examples

```
data(llq2022)
head(llq2022)
# Squared-error losses per model:
losses <- (llq2022[, c("AR1","AR22","AR22_Lasso","HAR","HARQ","ARFIMA")] -
             llq2022$rv)^2
colMeans(losses)
```

`llq2022_jnj` *Johnson & Johnson Realized-Variance Forecasts (Li, Liao & Quaedvlieg, 2022)*

Description

Daily 5-minute realized variance of Johnson & Johnson (NYSE: JNJ) with one-step-ahead out-of-sample forecasts from six competing models, plus one-day-lagged VIX as a conditioning variable. JNJ is the stock featured in Figures 2 and 3 of Li, Liao and Quaedvlieg (2022).

Usage

```
llq2022_jnj
```

Format

A data frame with the same columns as `llq2022`:

date Trading-day `Date`.

rv Realized variance (5-minute returns), percent-squared.

AR1, AR22, AR22_Lasso, HAR, HARQ, ARFIMA One-step-ahead forecasts of `rv`.

vix_lag VIX close from the previous trading day.

Source

Replication package of Li, Liao and Quaedvlieg (2022), https://zenodo.org/record/4884813. Realized variance series from Bollerslev, Patton and Quaedvlieg (2016). VIX from CBOE.

References

Bollerslev, T., Patton, A. J. and Quaedvlieg, R. (2016). Exploiting the errors: a simple approach for improved volatility forecasting. *Journal of Econometrics*, 192(1), 1-18.

Li, J., Liao, Z. and Quaedvlieg, R. (2022). Conditional Superior Predictive Ability. *Review of Economic Studies*, 89(2), 843-875.

`llq2022_uv_cspa` *Pairwise CSPA Rejection Counts Across 28 Stocks*

Description

Pre-computed rejection counts replicating Table_UV_CSPA.xlsx from the Li, Liao and Quaedvlieg (2022) replication package. For each of 28 stocks, pairwise conditional superior predictive ability (CSPA) tests are run for every benchmark-competitor pair across six realized-variance forecasting models; cell (k, l) counts the number of stocks for which the null "benchmark l conditionally dominates alternative k" is rejected at the 5% level. Conditioning variable is one-day-lagged VIX; loss is QLIKE.

Usage

```
llq2022_uv_cspa
```

Format

A list with the following components:

mine 6×6 integer matrix of rejection counts produced by `forecastdom::cspa_test` with `R = 10000` bootstrap reps and AIC pre-whitening, matching the call signature in `Empirics_Volatility.ox`.

paper 6×6 integer matrix from `Table_UV_CSPA.xlsx` in the LLQ replication package.

losses 28×6 matrix of mean QLIKE loss per stock and model.

tickers Character vector of the 28 tickers used.

models Character vector of the six model names.

level Significance level used (0.05).

R Number of bootstrap replications used.

Source

Computed by `data-raw/llq2022_uv_cspa.R` from the replication package of Li, Liao and Quaedvlieg (2022), https://zenodo.org/record/4884813.

References

Li, J., Liao, Z. and Quaedvlieg, R. (2022). Conditional Superior Predictive Ability. *Review of Economic Studies*, 89(2), 843-875.

`mse_f_test` *McCracken MSE-F Test for Equal Forecast Accuracy*

Description

Tests the null hypothesis of equal mean squared forecast error between the (restricted) benchmark and the (unrestricted) alternative nested model, using the MSE-F statistic of McCracken (2007). Under the null, the alternative does not reduce MSFE relative to the benchmark.

Usage

```
mse_f_test(e1, e2, h = 1L)
```

Arguments

`e1`	Numeric vector of forecast errors from the benchmark (restricted) model.
`e2`	Numeric vector of forecast errors from the alternative (unrestricted) model.
`h`	Integer; forecast horizon. Default `1`.

Details

The MSE-F statistic is:

$$\text{MSE-F} = (T - h + 1) \cdot \frac{\text{MSFE}_1 - \text{MSFE}_2}{\text{MSFE}_2}$$

where T is the number of out-of-sample observations and h is the forecast horizon.

Critical values are non-standard and depend on the number of extra regressors in the alternative model (k_2) and on $\pi = P/R$ (out-of-sample size relative to estimation window). See McCracken (2007, Table 1) for asymptotic critical values.

Value

A list with class `"mse_f_test"` containing:

`statistic`	The MSE-F test statistic.
`msfe1`	Mean squared forecast error of the benchmark.
`msfe2`	Mean squared forecast error of the alternative.
`n`	Number of out-of-sample observations.
`h`	Forecast horizon.

References

McCracken, M.W. (2007). Asymptotics for out of sample tests of Granger causality. *Journal of Econometrics*, 140(2), 719-752.

Clark, T.E. and McCracken, M.W. (2001). Tests of Equal Forecast Accuracy and Encompassing for Nested Models. *Journal of Econometrics*, 105(1), 85-110.

Examples

```
set.seed(42)
e1 <- rnorm(100)
e2 <- rnorm(100, sd = 0.9)
mse_f_test(e1, e2)
```

`nrtz2014` *Equity Premium and Technical Indicators (Neely, Rapach, Tu, & Zhou, 2014)*

Description

Monthly U.S. log equity premium (in percent) and the 14 binary technical-indicator predictors of Neely, Rapach, Tu, and Zhou (2014), "Forecasting the Equity Risk Premium: The Role of Technical Indicators." The macroeconomic predictors of the same paper are *not* included; for those, see rz2013.

Usage

```
nrtz2014
```

Format

A data frame with 733 rows and 16 variables. Sample period is 1950-12 to 2011-12 (monthly).

date First-of-month `Date`.

eq_prem Log equity premium, in percent: $100 \cdot [\log(1 + r_t) - \log(1 + rf_{t-1})]$.

MA_1_9, MA_1_12, MA_2_9, MA_2_12, MA_3_9, MA_3_12 Moving-average signals `MA(s,l)` for $s \in \{1, 2, 3\}$, $l \in \{9, 12\}$.

MOM_9, MOM_12 Momentum signals `MOM(m)` for $m \in \{9, 12\}$.

VOL_1_9, VOL_1_12, VOL_2_9, VOL_2_12, VOL_3_9, VOL_3_12 On-balance-volume signals `VOL(s,l)` for $s \in \{1, 2, 3\}$, $l \in \{9, 12\}$.

Details

Indicators are constructed from monthly S&P 500 closing prices and trading volume (Goyal-Welch updated dataset) following the rules in the paper:

- **MA(s,l)** – 1 if the s-month moving average of the price exceeds the l-month moving average, else 0.
- **MOM(m)** – 1 if today's price is at or above the price m months ago, else 0.
- **VOL(s,l)** – 1 if the s-month moving average of on-balance volume exceeds the l-month moving average, else 0.

Source

Replication archive of Neely et al. (2014): https://github.com/gabbocg/nrtz2014. The underlying price and volume data are from Amit Goyal's website (https://sites.google.com/view/agoyal145).

References

Neely, C. J., Rapach, D. E., Tu, J., and Zhou, G. (2014). Forecasting the equity risk premium: The role of technical indicators. *Management Science*, 60(7), 1772-1791.

Examples

```
data(nrtz2014)
head(nrtz2014)
# Fraction of months each technical indicator equals 1
colMeans(nrtz2014[, -(1:2)])
```

---

print.cspa_test *Print Method for CSPA Test Results*

Description

Displays a formatted summary of the CSPA test output, including the test statistic (theta), p-value, decision, and key estimation details.

Usage

```
## S3 method for class 'cspa_test'
print(x, digits = 4, ...)
```

Arguments

x	An object of class "cspa_test", as returned by cspa_test.
digits	Integer; number of decimal places for numeric output. Default 4.
...	Additional arguments (currently ignored).

Value

Invisibly returns x.

Examples

```
sim <- do_sim(J = 3, n = 250, a = 1, c = 0, rho_u = 0.4)
result <- cspa_test(sim$Y, sim$X, level = 0.05, trim = 2, R = 500L)
print(result)
```

qll_hat *Elliott-Muller Test for Time-Varying Coefficients*

Description

Computes the $q\hat{L}L$ statistic of Elliott and Muller (2006) for testing the null hypothesis that regression coefficients are constant over time against the alternative of general time variation.

Usage

```
qll_hat(y, X, Z = NULL, L = 0L)
```

Arguments

y	Numeric vector of length T; the dependent variable.
X	A T x k matrix of regressors linked to potentially time-varying coefficients.
Z	A T x d matrix of regressors with constant coefficients. Use NULL (default) if all coefficients may vary.
L	Integer; lag truncation for the Newey-West estimator of the variance. Default 0 (no correction).

Details

The test is based on optimal invariant statistics for the null of constant coefficients against local alternatives. The $q\hat{L}L$ statistic has non-standard critical values that depend on k; see Table 1 in Elliott and Muller (2006).

Selected critical values (5\

- k = 1: -5.91
- k = 2: -10.64
- k = 3: -15.78
- k = 4: -20.62
- k = 5: -25.87

Reject the null when $q\hat{L}L$ is below the critical value.

Value

A list with class "qll_hat" containing:

statistic	The $q\hat{L}L$ test statistic.
k	Number of potentially time-varying coefficients.
n	Number of observations.

References

Elliott, G. and Muller, U.K. (2006). Efficient Tests for General Persistent Time Variation in Regression Coefficients. *Review of Economic Studies*, 73(4), 907-940.

Examples

```
set.seed(42)
n <- 200
x <- matrix(rnorm(n * 2), n, 2)
y <- x %*% c(0.5, -0.3) + rnorm(n)
qll_hat(y, x)
```

`quaedvlieg2021` *Loss-Differential Path Forecasts from Quaedvlieg (2021)*

Description

Two example matrices of loss differentials at horizons $h = 1, \ldots, 20$ from the replication archive of Quaedvlieg (2021). `$uspa` is a configuration where average SPA holds but uniform SPA does not; `$aspa` is a configuration where both hold.

Usage

```
quaedvlieg2021
```

Format

A list with two elements, each a `T` x `20` numeric matrix:

uspa Loss differentials with aSPA but not uSPA.

aspa Loss differentials with both aSPA and uSPA.

Source

Author's replication archive distributed alongside Quaedvlieg (2021), *JBES* 39(1):40-53.

`rossi2006` *Bilateral Nominal Exchange Rates (Rossi, 2006)*

Description

Monthly nominal exchange rates against the U.S. dollar for five industrialised economies, used in the empirical exercise of Rossi (2006). Originally distributed as a Datastream extract in the replication archive at https://github.com/gabbocg/rossi2006.

Usage

```
rossi2006
```

Format

A data frame with 1,550 rows and 3 variables:

date First-of-month `Date`, from 1973-03-01 to 1998-12-01 (310 monthly observations per country).

country Factor with levels `"Canada"`, `"France"`, `"Germany"`, `"Italy"`, `"Japan"`.

fx Numeric. Nominal bilateral exchange rate level vs. the U.S. dollar, as supplied in the source file. The empirical exercise in Rossi (2006) works with `log(fx)` and its monthly differences.

Source

Replication archive of Rossi (2006): https://github.com/gabbocg/rossi2006.

References

Rossi, B. (2006). Are exchange rates really random walks? Some evidence robust to parameter instability. *Macroeconomic Dynamics*, 10(1), 20-38.

Meese, R. A. and Rogoff, K. (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics*, 14(1-2), 3-24.

Examples

```
data(rossi2006)
head(rossi2006)
with(subset(rossi2006, country == "Japan"),
     plot(date, log(fx), type = "l", main = "Japan log FX"))
```

`rrz2016` *Equity Premium and Short Interest Index (Rapach, Ringgenberg & Zhou, 2016)*

Description

Monthly U.S. log excess return on the S&P 500 and the short interest index (SII) used in Rapach, Ringgenberg, and Zhou (2016), "Short interest and aggregate stock returns" (*Journal of Financial Economics*, 121, 46-65). Sample period 1973-01 to 2014-12.

Usage

```
rrz2016
```

Format

A data frame with 504 rows and 3 variables:

date First-of-month `Date`.

r Log excess return on the S&P 500: $\log(1 + R_t) - \log(1 + rf_{t-1})$.

SII Short interest index, standardised linearly-detrended log EWSI.

Details

SII is constructed as standardised residuals from a linear regression of $\log(\text{EWSI}_t)$ on a time trend, where EWSI_t is the equal-weighted mean across all firms of the number of shares held short normalised by shares outstanding.

Source

Replication archive of Rapach, Ringgenberg, and Zhou (2016): https://github.com/gabbocg/rrz2016.

References

Rapach, D. E., Ringgenberg, M. C. and Zhou, G. (2016). Short interest and aggregate stock returns. *Journal of Financial Economics*, 121(1), 46-65.

Examples

```
data(rrz2016)
head(rrz2016)
plot(rrz2016$date, rrz2016$SII, type = "l",
     xlab = NULL, ylab = "SII")
```

`rz2013` *Equity Premium and Macro Predictors (Rapach & Zhou, 2013)*

Description

Monthly U.S. log equity premium and the 14 Goyal-Welch macroeconomic predictors used in Rapach and Zhou (2013), "Forecasting Stock Returns," Handbook of Economic Forecasting, Vol. 2A, Chapter 6. Sourced from the replication archive at https://github.com/gabbocg/rz2013.

Usage

```
rz2013
```

Format

A data frame with 1,009 rows and 16 variables. Sample period is 1926-12 to 2010-12 (monthly).

date First-of-month `Date`.

eq_prem Log equity premium: $\log(1+r_t) - \log(1+rf_{t-1})$, where r_t is the CRSP value-weighted S&P 500 return.

DP Log dividend-price ratio.

DY Log dividend yield.

EP Log earnings-price ratio.

DE Log dividend-payout ratio.

SVAR Stock variance: monthly sum of squared daily S&P 500 returns.

BM DJIA book-to-market ratio.

NTIS Net equity expansion.

TBL 3-month Treasury bill yield (decimal).

LTY Long-term government bond yield (decimal).

LTR Long-term government bond return (decimal).

TMS Term spread: `LTY - TBL`.

DFY Default yield spread: BAA minus AAA corporate yields.

DFR Default return spread: long-term corporate minus government bond return.

INFL_lag Inflation rate, lagged one month (decimal).

Source

Replication archive of Rapach and Zhou (2013): https://github.com/gabbocg/rz2013. The underlying data are from Amit Goyal's website (https://sites.google.com/view/agoyal145).

References

Rapach, D. E. and Zhou, G. (2013). Forecasting stock returns. In G. Elliott and A. Timmermann (Eds.), *Handbook of Economic Forecasting*, Vol. 2A, pp. 328-383. Elsevier.

Goyal, A. and Welch, I. (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4), 1455-1508.

Examples

```
data(rz2013)
head(rz2013)
summary(rz2013$eq_prem)
```

`spa_test` *Unconditional Superior Predictive Ability (USPA/SPA) Test*

Description

Tests the null hypothesis that a benchmark forecasting method has unconditional superior predictive ability over all competing alternatives, using Hansen's (2005) stationary bootstrap approach.

Usage

```
spa_test(Y, level, B = 5000L, q = 0.25)
```

Arguments

`Y`	An `n` x `J` matrix of loss differentials with respect to the benchmark. Positive values indicate the benchmark outperforms competitor `j` in period `t`.
`level`	Significance level (e.g., 0.05).
`B`	Integer; number of bootstrap replications. Default `5000`.
`q`	Numeric in $(0, 1)$; parameter for the geometric distribution in the stationary bootstrap. Controls the average block length $(1/q)$. Default `0.25` (average block length of 4).

Details

Let $d_{j,t} = -Y_{j,t}$, so that positive values mean competitor `j` beats the benchmark. The SPA test statistic is:

$$T^{SPA} = \max\left(\max_{1 \le j \le J} \frac{\sqrt{n}\,\bar{d}_j}{\hat{\omega}_j}, 0\right)$$

where $\bar{d}_j$ is the sample mean of $d_{j,t}$ and $\hat{\omega}_j^2$ is its Newey-West long-run variance.

The null hypothesis is rejected when the benchmark is not superior to all competitors, i.e., some competitor significantly outperforms it.

The bootstrap distribution uses Hansen's (2005) consistent recentring: competitors with $\sqrt{n}\,\bar{d}_j/\hat{\omega}_j < -\sqrt{2\log\log n}$ are treated as clearly inferior and recentred at zero instead of $\bar{d}_j$, so they do not inflate the critical value.

Value

A list with class "spa_test" containing:

statistic	The SPA test statistic (consistent).
pvalue	Bootstrap p-value.
reject	Logical; whether the null is rejected.
level	Significance level used.
d_bar	Sample mean of loss differentials for each competitor.
n	Number of observations.
J	Number of competitors.
B	Number of bootstrap replications.

References

Hansen, P.R. (2005). A Test for Superior Predictive Ability. *Journal of Business & Economic Statistics*, 23(4), 365-380.

White, H. (2000). A Reality Check for Data Snooping. *Econometrica*, 68(5), 1097-1126.

Examples

```
sim <- do_sim(J = 3, n = 250, a = 1, c = 0, rho_u = 0.4)
spa_test(sim$Y, level = 0.05, B = 500)
```

---

summary.cspa_test *Summary Method for CSPA Test Results*

Description

Provides a detailed summary of the CSPA test, including per-competitor diagnostics of the estimated conditional mean functions.

Usage

```
## S3 method for class 'cspa_test'
summary(object, digits = 4, ...)
```

Arguments

object	An object of class "cspa_test", as returned by cspa_test.
digits	Integer; number of decimal places. Default 4.
...	Additional arguments (currently ignored).

Value

Invisibly returns `object`.

Examples

```
sim <- do_sim(J = 3, n = 250, a = 1, c = 0, rho_u = 0.4)
result <- cspa_test(sim$Y, sim$X, level = 0.05, trim = 2, R = 500L)
summary(result)
```

`uspa_mh_test` *Uniform Multi-Horizon Superior Predictive Ability Test*

Description

Tests the null hypothesis that a benchmark forecasting method has *uniform* (i.e. simultaneous, at every horizon) superior predictive ability over a competitor in a multi-horizon setting, using Quaedvlieg's (2021) moving-block-bootstrap implementation of the uniform SPA test (`uSPA`).

Usage

```
uspa_mh_test(loss_diff, L, B = 999L, level = 0.05)
```

Arguments

`loss_diff`	A `T` x `H` matrix of loss differentials, defined as `L_bench - L_comp` at each horizon `h = 1, ..., H`. Positive values indicate the benchmark has higher loss (i.e. is worse) than the competitor at that horizon.
`L`	Integer; moving-block-bootstrap block length.
`B`	Integer; number of bootstrap replications. Default `999`.
`level`	Significance level (e.g. `0.05`). Default `0.05`.

Details

The uniform multi-horizon SPA test statistic is

$$t^{uSPA} = \min_{1 \le h \le H} \frac{\sqrt{T}\,\bar{d}_h}{\sqrt{\hat{\omega}_h}},$$

where $\bar{d}_h$ is the sample mean of the loss differential `L_bench - L_comp` at horizon `h`, and $\hat{\omega}_h$ is its Quadratic-Spectral HAC long-run variance estimate (Andrews, 1991).

Critical values are obtained from a moving-block bootstrap that recentres the loss differentials so the null is imposed (Quaedvlieg, 2021), with the "natural variance" estimator recomputed within each bootstrap replication. The resulting p-value is upper-tail: large positive values of the statistic (i.e. the benchmark has higher loss *uniformly* across horizons) lead to rejection. Equivalently,

rejection requires $\bar{d}_h > 0$ at every horizon; the test has limited power against alternatives where the benchmark is worse at only some horizons.

This is a port of the Matlab reference implementation accompanying Quaedvlieg (2021); the helpers `.mbb_indices`, `.mbb_variance`, and `.qs_lrvar` translate `Get_MBB_ID.m`, `MBB_Variance.m`, and `QS.m` respectively.

Value

A list with class `"uspa_mh_test"` containing:

`statistic`	The uSPA test statistic $t = \min_h \sqrt{T}\,\bar{d}_h / \sqrt{\hat{\omega}_h}$.
`pvalue`	Moving-block bootstrap p-value $B^{-1}\sum_b \mathbf{1}\{t < t_b^*\}$ (upper tail).
`reject`	Logical; whether the null is rejected at `level`.
`level`	Significance level used.
`L`	Block length used.
`B`	Number of bootstrap replications.
`T`	Number of time periods.
`H`	Number of forecast horizons.
`d_bar`	Sample mean of loss differentials at each horizon (length `H`).

References

Quaedvlieg, R. (2021). Multi-Horizon Forecast Comparison. *Journal of Business & Economic Statistics*, 39(1), 40-53.

Examples

```
set.seed(1)
ld <- matrix(rnorm(200 * 4, mean = 0.3), 200, 4)
uspa_mh_test(ld, L = 3, B = 199, level = 0.10)
```

---

wg2008 *Welch & Goyal (2008) Annual Equity-Premium Dataset*

Description

Annual equity premium and predictor series built from Welch and Goyal's original `PredictorData.xls` (annual sheet) — the data vintage shipped with the published 2008 paper. The equity premium is constructed as the log total return on the S&P 500 minus the log risk-free rate (per WG Section 1), and the log dividend-price ratio is $\log(D12_{t-1}) - \log(\text{Index}_{t-1})$.

Usage

```
wg2008
```

Format

A data frame with 134 rows and 13 columns covering 1872-2005:

year Calendar year.

Index S&P 500 price level.

D12 Twelve-month moving sum of dividends.

E12 Twelve-month moving sum of earnings.

Rfree Risk-free rate (decimal).

tbl Treasury-bill yield.

infl CPI inflation rate.

ntis Net equity expansion.

spret Total return on the S&P 500, $(P_t + D_t)/P_{t-1} - 1$.

logeqp Log equity premium, $\log(1 + \text{spret}_t) - \log(1 + \text{Rfree}_t)$.

log_dp_lag Lagged log dividend-price ratio, $\log(D12_{t-1}/\text{Index}_{t-1})$ — the predictor used in WG Table 1.

log_ep_lag Lagged log earnings-price ratio.

log_de_lag Lagged log dividend-payout ratio.

Source

Welch and Goyal (2008) original annual data file `PredictorData.xls`, distributed with the paper. A regularly updated version is maintained on Amit Goyal's website, https://sites.google.com/view/agoyal145.

References

Welch, I. and Goyal, A. (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4), 1455-1508.

Index

∗ **datasets**
  cm2001, 4
  gw2006, 13
  hl2005, 15
  llq2022, 17
  llq2022_jnj, 18
  llq2022_uv_cspa, 19
  nrtz2014, 21
  quaedvlieg2021, 24
  rossi2006, 25
  rrz2016, 26
  rz2013, 27
  wg2008, 31

aspa_mh_test, 3

cm2001, 4
csms, 5
cspa_test, 6, *8*, *22*, *29*
cspa_test_plot, 8
cw_test, 9

dm_test, 10
do_sim, 11

enc_new, 12

gw2006, 13
gw_test, 14

hl2005, 15

ivx_wald, 16

llq2022, 17, *18*
llq2022_jnj, 18
llq2022_uv_cspa, 19

mse_f_test, 20

nrtz2014, 21

print.cspa_test, 22

qll_hat, 23
quaedvlieg2021, 24

rossi2006, 25
rrz2016, 26
rz2013, *21*, 27

spa_test, *4*, 28
summary.cspa_test, 29

uspa_mh_test, *4*, 30

wg2008, 31