Package 'ForeComp'

Type Package

Title Size-Power Tradeoff Visualization for Equal Predictive Ability
of Two Forecasts

Version 1.0.0

Author Nathan Schor [aut],
Minchul Shin [aut, cre, cph]

Maintainer Minchul Shin <visiblehand@gmail.com>

Description Offers tools for visualizing and analyzing size and power properties
of tests for equal predictive accuracy, including Diebold-Mariano and
related procedures. Provides multiple Diebold-Mariano test implementations
based on fixed-smoothing approaches, including fixed-b methods such as
Kiefer and Vogelsang (2005) <doi:10.1017/S0266466605050565>, and
applications to tests for equal predictive accuracy as in Coroneo and
Iacone (2020) <doi:10.1002/jae.2756>, alongside conventional large-sample
approximations. HAR inference involves
nonparametric estimation of the long-run variance, and a key tuning
parameter (the truncation parameter) trades off size and power. Lazarus,
Lewis, and Stock (2021) <doi:10.3982/ECTA15404> theoretically characterize
the size-power frontier for the Gaussian multivariate location model.
'ForeComp' computes and visualizes the finite-sample size-power frontier of
the Diebold-Mariano test based on fixed-b asymptotics together with the
Bartlett kernel. To compute finite-sample size and power, it fits a best
approximating ARMA process to the input data and reports how the truncation
parameter performs and how robust testing outcomes are to its choice.

License GPL (>= 3)

Encoding UTF-8

URL https://github.com/mcmcs/ForeComp

BugReports https://github.com/mcmcs/ForeComp/issues

LazyData true

Depends R (>= 3.5.0), stats

Imports forecast, astsa, ggplot2, rlang

Suggests testthat (>= 3.0.0)

Config/testthat/edition 3

RoxygenNote 7.3.3

NeedsCompilation no

Repository CRAN

Date/Publication 2026-02-20 15:40:07 UTC

Contents

---

dm.test.bt   *Diebold-Mariano Test (Bartlett kernel, normal approximation)*

Description

Diebold-Mariano Test (Bartlett kernel, normal approximation). This is a two-sided test.

Usage

```
dm.test.bt(d, M = NA, Mopt = NA, cl = 0.05)
```

Arguments

d	loss differential
M	truncation parameter for the Bartlett kernel (if M = NA, then Mopt = 2 by default)
Mopt	option for automatic bandwidth, 1 if Lazarus et al. (2018), 2 if Newey and West (1994, default), 3 if Andrews textbook NW ($M = \lceil 0.75T^{1/3} \rceil$), 4 if CI baseline ($M = \lfloor T^{1/2} \rfloor$)
cl	confidence level (default = 0.05, i.e., 5%)

Value

This function returns a class with the following elements

- `rej` is a T/F value. TRUE (reject), FALSE (accept)
- `stat` is a test statistic
- `pval` is an associated p-value

Author(s)

Minchul Shin

| `dm.test.bt.fb` | *Diebold-Mariano Test (Bartlett kernel, fixed-b approximation)* |

Description

Diebold-Mariano Test (Bartlett kernel, fixed-b approximation). This is a two-sided test.

Usage

```
dm.test.bt.fb(d, M = NA, Mopt = NA, cl = 0.05)
```

Arguments

`d`	loss differential
`M`	truncation parameter for the Bartlett kernel (if `M = NA`, then `Mopt = 1` by default)
`Mopt`	option for automatic bandwidth, 1 if Lazarus et al. (2018, default), 2 if Newey and West (1994), 3 if Andrews textbook NW ($M = \lceil 0.75 T^{1/3} \rceil$), 4 if CI baseline ($M = \lfloor T^{1/2} \rfloor$)
`cl`	confidence level (default = 0.05, i.e., 5%), Only 0.05 (5%) or 0.10 (10%) are allowed.

Value

This function returns a class with the following elements

- `rej` is a T/F value. TRUE (reject), FALSE (accept)
- `stat` is a test statistic

Author(s)

Minchul Shin

---

dm.test.cnr.t *CNR (2017), |t|-test statitic*

Description

Randomization test based on asymptotic symmetry, |t|-test statistic

Usage

```
dm.test.cnr.t(d, q = 2, cl = 0.05)
```

Arguments

d	loss differential
q	number of blocks
cl	confidence level (default = 0.05, i.e., 5%)

Value

This function returns a class with the following elements

- `rej` is a T/F value. TRUE (reject), FALSE (accept)
- `stat` is a test statistic
- `pval` is an associated p-value

Author(s)

Minchul Shin

---

dm.test.cnr.w *CNR (2017), Wald-statistic*

Description

Randomization test based on asymptotic symmetry, Wald-statistic

Usage

```
dm.test.cnr.w(d, q = 2, cl = 0.05)
```

Arguments

d	loss differential
q	number of blocks
cl	confidence level (default = 0.05, i.e., 5%)

Value

This function returns a class with the following elements

- `rej` is a T/F value. TRUE (reject), FALSE (accept)

- `stat` is a test statistic

- `pval` is an associated p-value

Author(s)

Minchul Shin

---

dm.test.ewc.fb	*Diebold-Mariano Test (EWC, fixed-b approximation)*

Description

Diebold-Mariano Test (EWC, fixed-b approximation). This is a two-sided test.

Usage

```
dm.test.ewc.fb(d, B = NA, Bopt = NA, cl = 0.05)
```

Arguments

d	loss differential
B	truncation parameter for the EWC long-run variance estimator (if B = NA, then Bopt = 1 by default)
Bopt	option for automatic bandwidth, 1 if Lazarus et al. (2018, default, $B = \lfloor 0.4T^{2/3} \rfloor$)
cl	confidence level (default = 0.05, i.e., 5%)

Value

This function returns a class with the following elements

- `rej` is a T/F value. TRUE (reject), FALSE (accept)

- `stat` is a test statistic

- `pval` is an associated p-value

Author(s)

Minchul Shin

---

dm.test.im   *Ibragimov and Muller (2010)*

Description

t-Statistic based HAR-inference by Ibragimov and Muller (2010).

Usage

```
dm.test.im(d, q = 2, cl = 0.05)
```

Arguments

d	loss differential
q	number of blocks. When length(d) is not divisible by q, this implementation uses near-equal nonoverlapping blocks (block sizes differ by at most one) and issues a warning.
cl	confidence level (default = 0.05, i.e., 5%)

Value

This function returns a class with the following elements

- rej is a T/F value. TRUE (reject), FALSE (accept)

- stat is a test statistic

- pval is an associated p-value

Author(s)

Minchul Shin

---

dm.test.r   *Diebold-Mariano Test (with an original recommendation)*

Description

This function performs the Diebold-Mariano test with their original recommendation. Let d_t be a sequence of loss differential, $t = 1, 2, ..., T$. Then, the function performs a statistical test for the following null hypothesis:

$$E[d_t] = 0$$

This function follows the original recommendation by Diebold and Mariano (1995), where the long-run variance is estimated using the rectangular kernel truncated at $(h - 1)$. This function performs a two-sided test. Normal approximation is used to calculate critical values.

Usage

```
dm.test.r(d, h = 1, cl = 0.05)
```

Arguments

d	loss differential
h	h-step-ahead forecast (default = 1, i.e., comparing one-step-ahead forecasts)
cl	confidence level (default = 0.05, i.e., 5%)

Value

This function returns a class with the following elements

- rej is a T/F value. TRUE (reject), FALSE (accept)
- stat is a test statistic
- pval is an associated p-value

Author(s)

Minchul Shin

---

dm.test.r.m   *Diebold-Mariano Test (Modified-DM)*

Description

Diebold-Mariano Test (Modified-DM). Finite-sample modification to the original DM's test. This is a two-sided test.

Usage

```
dm.test.r.m(d, h = 1, cl = 0.05)
```

Arguments

d	loss differential
h	h-step-ahead forecast (default = 1, i.e., one-step-ahead forecasts)
cl	confidence level (default = 0.05, i.e., 5%)

Details

This function is based on dm.test in the "forecast" package on CRAN.

Value

This function returns a class with the following elements

- `rej` is a T/F value. TRUE (reject), FALSE (accept)
- `stat` is a test statistic
- `pval` is an associated p-value

Author(s)

Minchul Shin

---

dm.test.wpe.fb   *Diebold-Mariano Test (WPE, fixed-m approximation)*

Description

Diebold-Mariano Test (WPE, fixed-m approximation). This is a two-sided test. See Coroneo and Iacone (2020)

Usage

```
dm.test.wpe.fb(d, M = NA, Mopt = NA, cl = 0.05)
```

Arguments

d	loss differential
M	truncation parameter for the WPE long-run variance estimator (if M = NA, then Mopt = 1 by default)
Mopt	option for automatic bandwidth, 1 if Coroneo and Iacone's default value (M = floor(T^(1/3)))
cl	confidence level (default = 0.05, i.e., 5%)

Value

This function returns a class with the following elements

- `rej` is a T/F value. TRUE (reject), FALSE (accept)
- `stat` is a test statistic
- `pval` is an associated p-value

Author(s)

Minchul Shin

---

loss.diff.p	*Loss differential from two point forecasts*

Description

This function computes the loss differential from two point forecasts. Let $f_{1,t}$ and $f_{2,t}$ be point forecasts generated by the forecaster 1 (e.g., forecasting model 1) and the forecaster 2 (e.g., forecasting model 2), respectively. y_t is a forecasting target that both $f_{1,t}$ and $f_{2,t}$ try to forecast. For a given loss function the loss differential is defined and computed as

$$d_t = L(y_t, f_{1,t}) - L(y_t, f_{2,t})$$

for $t = 1, 2, ..., T$.

Usage

```
loss.diff.p(y, f1, f2, type = "quad", ...)
```

Arguments

y	actual value, y_t, for $t = 1, 2, ..., T$.
f1	point forecast 1, $f_{1,t}$, for $t = 1, 2, ..., T$.
f2	point forecast 2, $f_{2,t}$, for $t = 1, 2, ..., T$.
type	Type of loss functions to compare.

- "quad" : quadratic loss function (default)
- "abs" : absolute loss function
- "check" : check (quantile) loss function
- "linex" : linex loss function

...	Extra parameters to compute loss function.

- For type="quad" : no extra parameter is needed.
- For type="abs" : no extra parameter is needed.
- For type="check" : a scalar parameter (tau) is needed. tau $\in [0, 1]$. Default value is tau=0.5 if unspecified.
- For type="linex" : a scalar parameter (c) is needed. c is a real number. Default value is c=1.

Details

Type of loss functions. Define $e_t = y_t - f_t$. Quadratic loss function is defined as

$$L(y_t, f_t) = e_t^2$$

Absolute loss function is defined as

$$L(y_t, f_t) = |e_t|$$

Check loss function is defined as

$$L(y_t, f_t) = (\tau - 1(e_t < 0))e_t$$

Linex loss function is defined as

$$L(y_t, f_t) = \exp(ce_t) - ce_t - 1$$

If c is positive (negative) real number, then the linex function is exponential (linear) for e>0 and linear (exponential) for e<0.

A sample average of the returned sequence, $\bar{d}_t = \frac{1}{T}\sum_{t=1}^{T} d_t$, is the difference in average loss from forecaster 1 and forecaster 2. Positive $\bar{d}_t$ means that the forecaster 2 incurs lower expected loss compared to the forecaster 1. Negative $\bar{d}_t$ means that the forecaster 1 incurs lower expected loss compared to the forecaster 2. Statistical testing to test whether the loss differential is zero in expectation is possible. See, for example, dm.test.r.

Value

This function returns a sequence of the loss differential, d_t.

Author(s)

Minchul Shin

---

mikedata   *Simulated data 1*

Description

A dataset containing y, f1, f2 based on one of McCracken (2019)'s data generating processes. The name of the specification is "Unconditional-Rolling". R = 175, Rbar = 175, h = 12, P = 75; Note that the DGP is constructed in a way that there is no difference between f1 and f2 in terms of expected quadratic loss.

Usage

mikedata

Format

A data frame with 75 rows and 3 variables:

y Forecast target

f1 Point forecast based on forecasting model 1

f2 Point forecast based on forecasting model 2

Source

Simulated data.

PGDP	*Price Index for Gross National Product/Gross Domestic Product (PGDP)*

Description

Error Statistics for the Survey of Professional Forecasters for GNP/GDP Price Index

Usage

```
PGDP
```

Format

'PGDP' A data frame with 219 rows and 31 columns

Source

<https://www.philadelphiafed.org/-/media/frbp/assets/surveys-and-data/survey-of-professional-forecasters/data-files/pgdp/data_spf_error_statistics_pgdp_3_aic.xls?la=en&hash=148987D03D54DA5391A44F28CBC12729>

Plot_Tradeoff	*Visualizes the size distortion maximum power loss tradeoff from the Diebold-Mariano test for equal predictive accuracy*

Description

'Plot_Tradeoff' creates a plot to show sensitivity of statistical significance to the choice of bandwidth and how size distortion and maximum power loss vary. It is designed for the Diebold-Mariano test for equal predictive accuracy (Diebold and Mariano, 1995). For a size-power tradeoff plot, see Lazarus, Lewis, Stock, and Watson (2018) and Lazarus, Lewis, and Stock (2021).

Usage

```
Plot_Tradeoff(
  data,
  f1 = NULL,
  f2 = NULL,
  y = NULL,
  loss_function = NULL,
  n_sim = 1000,
  m_set = NULL,
  cl = 0.05,
  verbose = TRUE,
  no_m_label = FALSE,
  na_handling = c("zero", "drop", "error")
)
```

Arguments

`data`	A data frame.
`f1`	Column containing forecaster 1's predictions. Should be a string.
`f2`	Column containing forecaster 2's predictions. Should be a string.
`y`	Column containing the realized value for the outcome variable. Should be a string.
`loss_function`	The transformation applied to the forecast error. Defaults to squared error loss. The user supplied function should take two inputs and a scalar output, loss = loss_function(f, y). For example, quadratic loss function would be defined as `loss_function=function(f,y){(f-y)^2}`.
`n_sim`	The number of simulations used to generate the ARIMA model. Defaults to 1,000.
`m_set`	The truncation parameter. If `NULL`, the function constructs a default grid centered around the package default bandwidth for `dm.test.bt.fb`, $M_0 = \lceil 1.3\sqrt{T} \rceil$ (LLSW), where T is the effective sample size after applying `na_handling`. It should be a vector of integers with the values of M you would like to plot.
`cl`	Significance level used in the hypothesis test. Defaults to 0.05. Only 0.05 and 0.10 are currently supported because the fixed-b approximation is implemented for these levels.
`verbose`	TRUE to print out the progress to the console. Defaults to TRUE.
`no_m_label`	TRUE to plot without m labels. Defaults to FALSE.
`na_handling`	How to handle missing/non-finite loss differentials. `"zero"` (default) replaces them with 0 (Application 2 convention), `"drop"` removes them (Application 1 convention), and `"error"` stops if any are present.

Value

A list of length 2. The first element is a ggplot2 object of the size-power tradeoff. The second element is the underlying data used to construct the plot in element 1.

Author(s)

Nathan Schor and Minchul Shin

References

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

Lazarus, E., Lewis, D. J., Stock, J. H. & Watson, M. W. (2018), HAR Inference: Recommendations for Practice, *Journal of Business & Economic Statistics*, **36**(4), 541-559.

Lazarus, E., Lewis, D. J. & Stock, J. H. (2021), The Size-Power Tradeoff in HAR Inference, *Econometrica*, **89**(5), 2497-2516.

Examples

```
# A typical example
set.seed(1234)
output = Plot_Tradeoff(
  data = TBILL,
  f1   = "SPFfor_Step1",
  f2   = "NCfor_Step1",
  y    = "Realiz1",
  m_set = seq(from = 1, to = 70, by = 10)
)
output[[1]] # The first element is a ggplot2 object of the size-power tradeoff.
output[[2]] # The second element is the underlying data used to construct the plot in element 1.

# An example with a user supplied loss function
# To use the mean absolute error as a loss function rather than a quadratic loss function
set.seed(1234)
output = Plot_Tradeoff(
  data = TBILL,
  f1   = "SPFfor_Step1",
  f2   = "NCfor_Step1",
  y    = "Realiz1",
  loss_function = function(f,y){ abs(f-y) },
  m_set = seq(from = 1, to = 50, by = 10)
)

# An example without (f1, f2, y). The function will take the first three columns and use them
set.seed(1234)
tmpdata = TBILL[, c("SPFfor_Step1", "NCfor_Step1", "Realiz1")] # data with [f1, f2, y]
Plot_Tradeoff(
  data = tmpdata,
  m_set = seq(from = 1, to = 50, by = 10)
)
```

RGDP   *Real Gross National Product/Gross Domestic Product (RGDP)*

Description

Error Statistics for the Survey of Professional Forecasters for Real GNP/GDP

Usage

```
RGDP
```

Format

'RDGP' A data frame with 219 rows and 31 columns

Source

<https://www.philadelphiafed.org/-/media/frbp/assets/surveys-and-data/survey-of-professional-forecasters/data-files/rgdp/data_spf_error_statistics_rgdp_3_aic.xls?la=en&hash=3AC9E2D8A5299F93CA7E16CFAA974C22>

TBILL   *3-Month Treasury Bill Rate (TBILL)*

Description

Error Statistics for the Survey of Professional Forecasters for Treasury Bill Rate (Three Month)

Usage

    TBILL

Format

'TBILL' A data frame with 219 rows and 31 columns

Source

<https://www.philadelphiafed.org/-/media/frbp/assets/surveys-and-data/survey-of-professional-forecasters/data-files/tbill/data_spf_error_statistics_tbill_1_aic.xls?la=en&hash=F432350F84B0E4CCE9A1E9D399447CA9>

UNEMP   *Civilian Unemployment Rate (UNEMP)*

Description

Error Statistics for the Survey of Professional Forecasters for Unemployment Rate

Usage

    UNEMP

Format

'UNEMP' A data frame with 219 rows and 31 columns

Source

<https://www.philadelphiafed.org/-/media/frbp/assets/surveys-and-data/survey-of-professional-forecasters/data-files/unemp/data_spf_error_statistics_unemp_1_aic.xls?la=en&hash=4CAD0B11FEAB6C4D0F30C38965FE3354>

Index