

Package: mmbcv (via r-universe)

July 22, 2026

Title Multistate Model Bias-Corrected Robust Variance

Version 1.0.0

Description Computes robust and bias-corrected sandwich variance estimators for multi-state Cox models with clustered time-to-event data. Also provides Wald tests for heterogeneity, generalized least-squares linear trends, and order-restricted trends among transition-specific coefficients. The methodology extends the marginal Cox model bias-correction framework of Wang et al. (2023) <[doi:10.1002/bimj.202200113](https://doi.org/10.1002/bimj.202200113)> to the multi-state setting.

License MIT + file LICENSE

Encoding UTF-8

RoxygenNote 7.3.3

Suggests knitr, rmarkdown, survival, testthat (>= 3.0.0)

VignetteBuilder knitr

Config/testthat/edition 3

NeedsCompilation no

Depends R (>= 3.5.0)

Author Can Meng [aut, cre], Denise Esserman [aut], Fan Li [aut], Erich Greene [aut]

Maintainer Can Meng <can.meng@yale.edu>

Repository <https://cran.r-universe.dev>

Date/Publication 2026-07-22 06:46:10 UTC

RemoteUrl <https://github.com/cran/mmbcv>

RemoteRef HEAD

RemoteSha 6642676018d4c60f7f257a2911e9ac4e42c781b6

Contents

heterogeneity_test	2
linear_trend_test	4
MMBCV	8
msdat3	10
order_restricted_test	12
subset_by_transition	16

Index	17
--------------	-----------

heterogeneity_test	<i>Omnibus Wald test for equality of correlated coefficients</i>
--------------------	--

Description

Performs a non-directional omnibus Wald test of whether a selected collection of correlated coefficient parameters is equal. The null hypothesis states that all selected parameters are equal; the unrestricted alternative states that at least one parameter differs.

The test can be applied to any collection of coefficient estimates for which a full variance-covariance matrix is available. No natural ordering or numeric scores are required.

Usage

```
heterogeneity_test(beta, vcov, index = NULL)
```

Arguments

beta	A finite numeric vector containing the coefficient estimates. At least two coefficients must be selected for testing. If index = NULL, all coefficients in beta are tested.
vcov	A square numeric variance-covariance matrix corresponding to beta. Its dimensions must be length(beta) by length(beta), and it must include the off-diagonal covariances between coefficient estimates. The covariance matrix of the selected contrasts must be nonsingular.
index	Optional integer, logical, or character vector selecting the coefficients to test. If NULL, all coefficients in beta are tested. The same selection is automatically applied to the rows and columns of vcov. Users may instead supply an already-subset coefficient vector and covariance matrix. The supplied order determines the ordering of the returned adjacent contrasts.

Details

Let $\hat{\beta}$ denote the vector of selected coefficient estimates, and let $\hat{V}$ denote their variance-covariance matrix. The adjacent contrast matrix D defines

$$D\hat{\beta} = (\hat{\beta}_2 - \hat{\beta}_1, \dots, \hat{\beta}_R - \hat{\beta}_{R-1})^T.$$

The null and alternative hypotheses are

$$H_0 : D\beta = 0,$$

which is equivalent to

$$H_0 : \beta_1 = \beta_2 = \dots = \beta_R,$$

against

$$H_A : D\beta \neq 0.$$

The Wald statistic is

$$W = (D\hat{\beta})^T (D\hat{V}D^T)^{-1} (D\hat{\beta}).$$

Under the null hypothesis, W has an asymptotic chi-square distribution with $R - 1$ degrees of freedom.

This is a non-directional omnibus test. Although its p-value is calculated from the upper tail of the chi-square distribution, it is not a directional one-sided test. The quadratic statistic is unchanged when $D\hat{\beta}$ is replaced by $-D\hat{\beta}$, so departures from equality in every direction are included in the alternative.

When two coefficients are selected, the test is equivalent to a two-sided Wald Z test because $W = Z^2$. When more than two coefficients are selected, it is more precisely described as a non-directional omnibus test. Directional linear or monotone alternatives should be evaluated using [linear_trend_test\(\)](#) or [order_restricted_test\(\)](#), respectively.

Value

An object of class "htest" containing:

statistic The omnibus Wald statistic W .

parameter The chi-square degrees of freedom, equal to $R - 1$.

p.value The upper-tail chi-square p-value.

method A character string identifying the statistical test.

alternative A character string describing the unrestricted alternative hypothesis.

coefficients The selected coefficient estimates $\hat{\beta}$.

vcov The variance-covariance matrix corresponding to the selected coefficients.

index The resolved positions of the selected coefficients in the original coefficient vector.

differences The estimated adjacent contrasts $D\hat{\beta}$.

contrast_matrix The adjacent contrast matrix D .

contrast_vcov The variance-covariance matrix of the adjacent contrasts, $D\hat{V}D^T$.

The "htest" class provides the standard printed format for hypothesis tests in R. Additional returned components can be accessed using \$.

References

- Wald, A. (1943). Tests of statistical hypotheses concerning several parameters when the number of observations is large. *Transactions of the American Mathematical Society*, 54, 426–482.
- Lehmann, E. L., and Romano, J. P. (2022). *Testing Statistical Hypotheses*, 4th ed. Springer, Cham.
- Cox, D. R., and Hinkley, D. V. (1974). *Theoretical Statistics*. Chapman and Hall, London.

See Also

`linear_trend_test()`, `order_restricted_test()`

Examples

```
beta <- c(-0.3, -0.2, -0.1, 0)
V <- diag(0.01, 4)
heterogeneity_test(beta, V)
```

linear_trend_test	<i>Linear trend test for a sequence of correlated coefficients</i>
-------------------	--

Description

Tests for a linear relationship between a selected sequence of correlated coefficient parameters and their assigned numeric scores. Although intended primarily for transition-specific treatment effects, the procedure can be applied to any ordered coefficient sequence with a full variance-covariance matrix.

Usage

```
linear_trend_test(
  beta,
  vcov,
  index = NULL,
  scores = NULL,
  alternative = c("greater", "less", "two.sided")
)
```

Arguments

beta	A finite numeric vector of coefficient estimates. If index is NULL, beta should contain only the coefficients to be tested, arranged in order for the trend model. At least two coefficients are required.
vcov	A square numeric variance-covariance matrix corresponding to beta. Its dimensions must be <code>length(beta)</code> by <code>length(beta)</code> . The full covariance matrix, including off-diagonal covariances, is required because the generalized least-squares calculation accounts for correlation among the coefficient estimates.
index	An optional integer, logical, or character vector selecting the coefficients to include in the trend test. If NULL, all coefficients in beta are used. The same selection is automatically applied to the rows and columns of vcov. Character selection requires coefficient names. The supplied order defines the order used in the trend analysis. Users may instead supply an already-subset coefficient vector and covariance matrix.

scores	A finite numeric vector assigning an ordered position s_r to each coefficient included in the linear trend model. The selected scores must be strictly increasing. If omitted, equally spaced scores $0, 1, \dots, R - 1$ are used, where R is the number of selected coefficients. Scores should be supplied when the selected coefficients correspond to unequally spaced ordered positions. They are also required when <code>index</code> selects nonconsecutive coefficient positions, so that the intended spacing is specified explicitly. <code>scores</code> must either have length equal to the number of selected coefficients or have length <code>length(beta)</code> . In the latter case, the scores corresponding to the selected coefficients are extracted using <code>index</code> . Consequently, the score vector used in the trend model always has length equal to the number of selected coefficients.
alternative	A character string specifying the alternative hypothesis: "greater" tests for a positive trend slope, "less" tests for a negative trend slope, and "two.sided" tests for any nonzero trend slope. Partial matching is allowed.

Details

Suppose the selected coefficient parameters are treated as an ordered sequence $\hat{\beta} = (\hat{\beta}_1, \dots, \hat{\beta}_R)^T$, with estimated variance-covariance matrix $\hat{V}$. The test represents the coefficients using the linear trend model

$$\beta_r = \eta + \gamma s_r, \quad r = 1, \dots, R,$$

where η is an intercept, γ is the trend slope, and s_r is the strictly increasing numeric score assigned to ordered-position r . The scores define the ordering and relative spacing used in the linear trend model. They may represent ordered categories, measurement occasions, dose levels, transition positions, or any other meaningful ordered scale.

Let

$$A = \begin{pmatrix} 1 & s_1 \\ \vdots & \vdots \\ 1 & s_R \end{pmatrix}$$

be the design matrix. The generalized least-squares estimator of $\theta = (\eta, \gamma)^T$ is

$$\hat{\theta} = (A^T \hat{V}^{-1} A)^{-1} A^T \hat{V}^{-1} \hat{\beta}.$$

Its estimated variance-covariance matrix is

$$\widehat{\text{Var}}(\hat{\theta}) = (A^T \hat{V}^{-1} A)^{-1}.$$

The test statistic for the trend slope is

$$Z = \frac{\hat{\gamma}}{\text{SE}(\hat{\gamma})}.$$

Under the boundary point of the null parameter space $\gamma = 0$, Z has an asymptotic standard-normal distribution.

The hypotheses depend on `alternative`:

"greater" $H_0 : \gamma \leq 0$ against $H_A : \gamma > 0$. The p-value is $1 - \Phi(Z)$.

"less" $H_0 : \gamma \geq 0$ against $H_A : \gamma < 0$. The p-value is $\Phi(Z)$.

"two.sided" $H_0 : \gamma = 0$ against $H_A : \gamma \neq 0$. The p-value is $2\{1 - \Phi(|Z|)\}$.

Here, Φ denotes the standard-normal cumulative distribution function.

The direction of the trend is defined relative to increasing score values. A positive slope indicates that the coefficients tend to increase as the assigned scores increase, whereas a negative slope indicates that they tend to decrease. The appropriate alternative should be selected according to the coefficient scale and the scientific question.

The spacing of the assigned scores affects the fitted trend. With the default scores, $\hat{\gamma}$ represents the estimated change in the coefficient per one unit increase in ordered position. With user-specified scores, it represents the estimated change per one score unit. Replacing the scores by a positive affine transformation changes the numerical scale of the slope but not its Wald statistic or p-value. Changing the relative spacing of the scores can change the fitted trend and the test result.

This test assesses whether the selected coefficients have a nonzero linear trend with respect to their assigned scores. A significant result indicates a nonzero generalized least-squares slope; it does not establish that a linear model provides an exact or adequate description of the coefficient pattern. The observed coefficient estimates are not required to be monotone. Use `heterogeneity_test()` to test for any departure from equality without imposing a trend shape. Use `order_restricted_test()` to test a monotone increasing or decreasing alternative without assuming linearity or equal-sized changes between adjacent coefficients.

Value

An object of class "htest" containing:

statistic The standard-normal Wald statistic Z .

p.value The p-value corresponding to the selected alternative hypothesis.

null.value The null value of the trend slope, equal to zero.

alternative The selected alternative: "greater", "less", or "two.sided".

method A character string identifying the generalized least-squares linear-trend test.

slope The generalized least-squares estimate $\hat{\gamma}$ of the trend slope.

stderr The estimated standard error of $\hat{\gamma}$.

coefficients The selected coefficient estimates.

vcov The variance-covariance matrix corresponding to the selected coefficients.

index The resolved positions of the selected coefficients in the original coefficient vector.

scores The resolved ordered-position scores used as predictor values in the linear trend model fitted by generalized least squares.

intercept The generalized least-squares estimate $\hat{\eta}$ of the intercept.

coefficient_covariance The estimated 2×2 variance-covariance matrix of the intercept and trend-slope estimates.

The "htest" class provides the standard printed format for hypothesis tests in R. Additional returned components can be accessed using \$, for example, `result$slope`, `result$stderr`, and `result$scores`.

References

- Cochran, W. G. (1954). Some methods for strengthening the common chi-squared tests. *Biometrics*, 10, 417–451.
- Armitage, P. (1955). Tests for linear trends in proportions and frequencies. *Biometrics*, 11, 375–386.
- Aitken, A. C. (1935). On least squares and linear combinations of observations. *Proceedings of the Royal Society of Edinburgh*, 55, 42–48.
- Seber, G. A. F., and Lee, A. J. (2003). *Linear Regression Analysis*, 2nd ed. John Wiley & Sons, Hoboken, NJ.
- Kariya, T., and Kurata, H. (2004). *Generalized Least Squares*. John Wiley & Sons.

See Also

[heterogeneity_test](#), [order_restricted_test](#)

Examples

```
beta <- c(
  beta1 = -0.40,
  beta2 = -0.30,
  beta3 = -0.20,
  beta4 = -0.10,
  beta5 = 0.00
)
V <- diag(0.01, 5)
dimnames(V) <- list(names(beta), names(beta))

# Increasing linear trend with equally spaced transitions
result <- linear_trend_test(
  beta,
  V,
  alternative = "greater"
)

result$slope
result$stderr

# Preserve the original spacing after omitting beta2
keep <- c(1, 3, 4, 5)
linear_trend_test(
  beta,
  V,
  index = keep,
  scores = 0:4,
  alternative = "greater"
)

# Unequally spaced transition scores
linear_trend_test(
  beta,
```

```
V,
scores = c(0, 1, 3, 6, 10),
alternative = "two.sided"
)
```

MMBCV

MMBCV: Bias-corrected cluster-robust sandwich variance for multi-state Cox models

Description

Computes several sandwich variance estimators (robust, MR, MD, FG, KC, MBN and combinations) for multistate Cox models using cluster-level score aggregation, with Breslow or Efron ties.

Usage

```
MMBCV(
  fit,
  data,
  StartTime,
  StopTime,
  ClusterID,
  SubjectID,
  Event,
  tie = c("breslow", "efron"),
  details = FALSE,
  detail_names = NULL,
  kc_tol = 1e-10
)
```

Arguments

<code>fit</code>	A fitted multistate Cox model object with components <code>cmap</code> , <code>rmap</code> , <code>states</code> , and <code>coefficients</code> . Typically produced by workflows in the survival package.
<code>data</code>	A <code>data.frame</code> in counting-process form that matches the fit.
<code>StartTime</code> , <code>StopTime</code>	Unquoted column names for start/stop times.
<code>ClusterID</code>	Unquoted column name for the cluster identifier.
<code>SubjectID</code>	Unquoted column name for the subject identifier.
<code>Event</code>	Unquoted column name for the event/state indicator used to define transition-specific events.
<code>tie</code>	Tie-handling method: "breslow" or "efron".
<code>details</code>	Logical; if TRUE, returns additional intermediate quantities.
<code>detail_names</code>	Optional character vector specifying which extra components of the full sandwich object to append when <code>details = TRUE</code> .
<code>kc_tol</code>	Numeric tolerance used in the KC positive-semidefinite adjustment.

Value

A list containing variance-covariance matrices:

robust Standard cluster-robust sandwich variance estimator.

varMR Martingale-residual (MR) bias-corrected variance estimator.

varMD Mancini-DeRouen (MD) bias-corrected variance estimator.

varMDMR Hybrid Mancini-DeRouen plus martingale-residual (MDMR) bias-corrected variance estimator.

varFG Fay-Graubard (FG) bias-corrected variance estimator.

varFGMR Hybrid Fay-Graubard plus martingale-residual (FGMR) bias-corrected variance estimator.

varKC Kauermann-Carroll (KC) bias-corrected variance estimator.

varKCMR Hybrid Kauermann-Carroll plus martingale-residual (KCMR) bias-corrected variance estimator.

varMBN Morel-Bokossa-Neerchal (MBN) bias-corrected variance estimator.

varMBNMR Hybrid Morel-Bokossa-Neerchal plus martingale-residual (MBNMR) bias-corrected variance estimator.

If `details = TRUE`, additional intermediate matrices may also be returned.

References

Wang, X., Turner, E. L., and Li, F. (2023). Improving sandwich variance estimation for marginal Cox analysis of cluster randomized trials. *Biometrical Journal*, 65, 2200113.

Schaubel, D. E. (2005). Adjusting analytic inference for small samples in recurrent event regression. *Statistics in Medicine*, 24(19), 3037–3051.

Kauermann, G., and Carroll, R. J. (2001). A note on the efficiency of sandwich covariance matrix estimation. *Journal of the American Statistical Association*, 96(456), 1387–1396.

Fay, M. P., and Graubard, B. I. (2001). Small-sample adjustments for Wald-type tests using sandwich estimators. *Biometrics*, 57(4), 1198–1206.

Mancini, L. A., and DeRouen, T. A. (2001). A covariance estimator for GEE with improved small-sample properties. *Biometrics*, 57(1), 126–134.

Morel, J. G., Bokossa, M., and Neerchal, N. K. (2003). Small sample correction for the variance of GEE estimators. *Biometrical Journal*, 45(4), 395–409.

Therneau, T., Crowson, C., and Atkinson, E. (2026). *Multi-state models and competing risks*. Vignette for the **survival** package. See `vignette("compete", package = "survival")` or <https://CRAN.R-project.org/package=survival>.

See Also

[coxph](#), [Surv](#)

Examples

```
data("msdat3")

# Example requires the suggested package 'survival'
if (requireNamespace("survival", quietly = TRUE)) {
  fit <- survival::coxph(
    list(
      survival::Surv(Tstart, Tstop, event) ~ 1,
      state("(s0)":state("S1") + state("S1"):state("S2") + state("S2"):state("S3")) ~
        Z + X,
      state("(s0)":state("D") + state("S1"):state("D") + state("S2"):state("D") +
        state("S3"):state("D")) ~ (Z + X)/common
    ),
    data = msdat3, id = id, ties = "breslow",
    timefix = FALSE
  )

  MMBCV(
    fit      = fit,
    data     = msdat3,
    StartTime = Tstart,
    StopTime  = Tstop,
    ClusterID = clus_id,
    SubjectID = id,
    Event     = event,
    tie       = "breslow",
    details   = FALSE)
}
```

msdat3

Clustered multistate simulated dataset (counting-process format)

Description

A simulated clustered, subject-level multistate dataset in long (counting-process) format. Treatment is assigned at the cluster level and baseline covariates are available at both the cluster and individual levels. Each subject may contribute multiple rows, corresponding to transitions between states over time.

Usage

```
data(msdat3)
```

Format

A data frame with 3,004 rows and 9 variables:

id Subject identifier. There are 1,500 unique subjects.

Z Cluster-level treatment assignment indicator. $Z = 1$ indicates the cluster is assigned to treatment; $Z = 0$ indicates control.

X Individual-level baseline covariate generated from a standard normal distribution.

clus_id Cluster identifier. There are 30 clusters.

from Starting state for the current observation interval (e.g., (s0)).

to Ending state for the current observation interval (e.g., S1, S2, S3, D, or censor).

Tstart Start time of the interval.

Tstop Stop time of the interval.

event End state at Tstop. Values include censor, S1, S2, S3, and D, where D is the absorbing state.

Details

The multistate process uses the following state labels:

- (s0): initial (starting) state.
- S1: first intermediate transition state.
- S2: second intermediate transition state.
- S3: third intermediate transition state.
- D: absorbing (terminal) state.

Each row represents a time interval $[Tstart, Tstop]$ during which the subject is at risk for the transition indicated by **from** and **to**. The variable **event** records the state at which the interval ends. **event** = "censor" indicates the observation is right-censored at Tstop (i.e., no transition is observed at the end of the interval).

The distribution of **event** in this dataset is:

- censor: 1050
- S1: 837
- S2: 450
- S3: 217
- D: 450

Source

Simulated data bundled with the package for examples and testing.

Examples

```
data("msdat3")
nrow(msdat3)
length(unique(msdat3$id))
length(unique(msdat3$clus_id))
table(msdat3$event)
```

order_restricted_test *Order-restricted test of equality against a monotone alternative*

Description

Tests whether a selected sequence of correlated coefficient parameters is constant against the alternative that the parameters follow a monotone increasing or decreasing order. The procedure uses a covariance-weighted projection onto an order-restricted cone and a chi-bar-square reference distribution.

Although the function is designed primarily for comparing transition-specific treatment effects, it can be applied to any ordered vector of coefficient estimates for which a full variance-covariance matrix is available.

Usage

```
order_restricted_test(
  beta,
  vcov,
  index = NULL,
  alternative = c("increasing", "decreasing"),
  weight_method = c("auto", "simulation"),
  nsim = 100000L,
  seed = NULL,
  weights = NULL,
  tol = 1e-10
)
```

Arguments

beta	A finite numeric vector containing the coefficient estimates. At least two coefficients must be selected for testing. The positions of the coefficients, or the order specified through index, define the sequence to which the monotonicity restriction is applied. Optional coefficient names are retained in the returned object.
vcov	A square numeric variance-covariance matrix corresponding to beta. Its dimensions must be length(beta) by length(beta), and it must include the off-diagonal covariances between coefficient estimates. If beta and vcov are named, their names must agree. The covariance matrix of the selected adjacent contrasts must be positive definite.
index	An optional integer, logical, or character vector selecting the coefficients to be tested. If index = NULL, all coefficients are used in their original order. An integer vector supplies coefficient positions, a logical vector must have length length(beta), and a character vector selects coefficients by name. Character selection requires names for beta or matching dimension names for vcov. The order of the resolved positions defines the order restriction. At least two distinct coefficients must be selected.

alternative	The direction of the order-restricted alternative. "increasing" tests for $\beta_1 \leq \beta_2 \leq \dots \leq \beta_R$, with at least one strict inequality. "decreasing" tests the reversed ordering $\beta_1 \geq \beta_2 \geq \dots \geq \beta_R$.
weight_method	The method used to calculate the chi-bar-square mixture weights. Available choices are: • "auto": use analytic weights when there are one or two order constraints and simulation otherwise; • "simulation": estimate the weights by simulation regardless of whether analytic weights are available. This argument is ignored when weights is supplied.
nsim	A positive integer giving the number of simulation draws used to estimate the chi-bar-square weights. It is used only when weights are calculated by simulation. Larger values reduce Monte Carlo error but increase computation time. The default is 100000.
seed	An optional non-negative integer used as the random-number seed when the chi-bar-square weights are estimated by simulation. Supplying a seed makes the estimated weights and p-value reproducible. It has no effect when analytic or user-supplied weights are used.
weights	An optional numeric vector of precomputed chi-bar-square weights. If R coefficients are selected, the vector must contain R non-negative values corresponding, in order, to degrees of freedom $0, 1, \dots, R - 1$. The weights must sum to one, within numerical tolerance. When supplied, weight calculation by analytic formulas or simulation is bypassed.
tol	A positive finite numeric tolerance used in covariance-matrix checks, covariance-weighted projection, determination of active cone faces, and validation of supplied weights. The default is $1e-10$. This is primarily an advanced numerical-control argument.

Details

Suppose R coefficient parameters are selected and collected in

$$\beta = (\beta_1, \dots, \beta_R)^T.$$

The null hypothesis states that all selected parameters are equal:

$$H_0 : \beta_1 = \beta_2 = \dots = \beta_R.$$

For an increasing alternative, the alternative hypothesis is

$$H_A : \beta_1 \leq \beta_2 \leq \dots \leq \beta_R,$$

with at least one strict inequality. For a decreasing alternative, all inequalities are reversed. Thus, the test is directional. It differs from `heterogeneity_test()`, whose alternative includes departures from equality in any direction.

Let D denote the directional adjacent contrast matrix. For an increasing alternative,

$$\hat{\delta} = D\hat{\beta} = (\hat{\beta}_2 - \hat{\beta}_1, \dots, \hat{\beta}_R - \hat{\beta}_{R-1})^T.$$

For a decreasing alternative, the sign of D is reversed, so that positive contrasts always represent movement in the direction specified by alternative. If $\hat{V}$ is the variance-covariance matrix of $\hat{\beta}$, then

$$\hat{\Omega} = D\hat{V}D^T$$

is the variance-covariance matrix of the directional adjacent contrasts. The hypotheses can therefore be written as

$$H_0 : \delta = 0 \quad \text{versus} \quad H_A : \delta \geq 0, \delta \neq 0.$$

The observed contrast vector is projected onto the nonnegative orthant using the metric induced by $\hat{\Omega}^{-1}$:

$$\tilde{\delta} = \underset{\theta \geq 0}{\operatorname{argmin}} (\hat{\delta} - \theta)^T \hat{\Omega}^{-1} (\hat{\delta} - \theta).$$

The order-restricted statistic is

$$T = \hat{\delta}^T \hat{\Omega}^{-1} \hat{\delta} - (\hat{\delta} - \tilde{\delta})^T \hat{\Omega}^{-1} (\hat{\delta} - \tilde{\delta}).$$

The statistic measures the improvement in covariance-weighted fit obtained by moving from the equality-constrained null point to the closest point in the monotone alternative cone. It is non-negative. If the observed contrast vector already lies in the alternative cone, its projection is unchanged and the statistic can equal the ordinary omnibus Wald statistic. The two procedures nevertheless have different reference distributions and generally different p-values.

Under the null hypothesis, the statistic has an asymptotic chi-bar-square distribution:

$$\bar{\chi}^2 = \sum_{k=0}^{R-1} w_k \chi_k^2,$$

where the non-negative weights w_k sum to one and depend on the covariance structure and geometry of the order cone. The p-value is obtained by combining the upper-tail probabilities of the chi-square components using these weights. If the statistic is numerically zero, the p-value is set to one.

With one or two order constraints, corresponding to two or three selected coefficients, analytic weights are available. For larger problems, `weight_method = "auto"` estimates the weights by simulation. The current simulation algorithm supports at most 12 order constraints; for a larger problem, precomputed weights must be supplied through `weights`.

Only the ordering of the selected coefficients is relevant. Numeric distances between transitions or other ordered categories are not required. For example, `index = c(1, 3, 4, 5)` tests the order of $(\beta_1, \beta_3, \beta_4, \beta_5)^T$. Unlike `linear_trend_test()`, no transition scores are needed because the order-restricted test does not assume a constant linear change per unit of transition score.

The function reports a statistic and p-value but does not require a significance level. A user-selected level α can be applied by comparing `result$p.value < alpha`.

Value

An object of class "htest" containing:

statistic The non-negative order-restricted test statistic, named "Tord".

p.value The upper-tail p-value calculated from the chi-bar-square reference distribution.

method A character string identifying the testing procedure and its chi-bar-square reference distribution.

alternative A character string describing the increasing or decreasing monotone alternative hypothesis.

coefficients The selected coefficient estimates in the order used by the test.

vcov The variance-covariance matrix of the selected coefficients.

index The resolved positions of the selected coefficients in the original beta vector.

differences The estimated directional adjacent contrasts. For an increasing alternative, these are later-minus-earlier differences $\hat{\beta}_{r+1} - \hat{\beta}_r$. For a decreasing alternative, their signs are reversed.

projected_differences The covariance-weighted projection of differences onto the nonnegative orthant representing the monotone alternative.

contrast_matrix The directional adjacent contrast matrix D used to calculate differences.

contrast_vcov The variance-covariance matrix $D\hat{V}D^T$ of the directional adjacent contrasts.

weights The chi-bar-square mixture weights corresponding to degrees of freedom $0, 1, \dots, R - 1$.

df The degrees of freedom associated with the chi-square mixture components.

weight_method The method actually used to obtain the weights: analytic calculation, projection simulation, or user-supplied weights.

nsim The number of simulation draws actually used. This is zero when analytic or user-supplied weights are used.

weight_mcse The estimated Monte Carlo standard error of each simulated chi-bar-square weight. These values are zero for analytic or user-supplied weights.

p.value.mcse The estimated Monte Carlo standard error of the p-value arising from simulation of the chi-bar-square weights. It is zero for analytic or user-supplied weights.

projection_method A character string identifying the numerical method used for the covariance-weighted cone projection.

The "htest" class provides R's standard printed format for hypothesis tests. Additional components not shown in the printed output can be accessed with \$.

References

- Robertson, T., Wright, F. T., and Dykstra, R. L. (1988). *Order Restricted Statistical Inference*. John Wiley & Sons, New York.
- Kudo, A. (1963). A multivariate analogue of the one-sided test. *Biometrika*, 50, 403–418.
- Shapiro, A. (1985). Asymptotic distribution of test statistics in the analysis of moment structures under inequality constraints. *Biometrika*, 72, 133–144.
- Silvapulle, M. J., and Sen, P. K. (2005). *Constrained Statistical Inference: Inequality, Order, and Shape Restrictions*. John Wiley & Sons, Hoboken, NJ.

See Also

[heterogeneity_test\(\)](#), [linear_trend_test\(\)](#)

Examples

```
beta <- c(beta1 = -0.30, beta2 = -0.20,
          beta3 = -0.10, beta4 = 0)
V <- diag(0.01, 4)

result <- order_restricted_test(
  beta,
  V,
  alternative = "increasing",
  nsim = 2000,
  seed = 47
)

result$contrast_matrix
result$weights
```

subset_by_transition *Subset a dataset by transition using a row map*

Description

Subset a dataset by transition using a row map

Usage

```
subset_by_transition(dataA, dataB, transition)
```

Arguments

dataA	Full dataset to subset.
dataB	Two-column object: column 1 = row indices in dataA, column 2 = transition id.
transition	Character vector of transition names used as output list names.

Value

A named list of data frames, one per transition.

Index

* **datasets**

msdat3, [10](#)

coxph, [9](#)

heterogeneity_test, [2](#), [7](#)

heterogeneity_test(), [6](#), [16](#)

linear_trend_test, [4](#)

linear_trend_test(), [3](#), [4](#), [16](#)

MMBCV, [8](#)

msdat3, [10](#)

order_restricted_test, [7](#), [12](#)

order_restricted_test(), [3](#), [4](#), [6](#)

subset_by_transition, [16](#)

Surv, [9](#)