

Package ‘posc’

October 21, 2026

Type Package

Title Probability of Outcome Superiority Curves

Version 0.2.0

Description Computes, plots and compares probability of outcome superiority curves (POSCs). For two people who differ by a given amount on a continuous predictor, a POSC gives the probability that the person with the higher predictor score also has the higher outcome, a probabilistic and decision-oriented alternative to reporting a correlation coefficient. Curves can be computed from a correlation under bivariate normality, from a correlation matrix for an optimally weighted composite of predictors, from a fitted linear model, or estimated directly from raw data with a flexible spline model and bootstrap confidence bands.

License MIT + file LICENSE

URL <https://matthewbjane.com/posc/>,
<https://github.com/MatthewBJane/posc>

BugReports <https://github.com/MatthewBJane/posc/issues>

Depends R (>= 4.1.0)

Imports ggplot2 (>= 3.4.0), splines, stats, utils

Suggests testthat (>= 3.0.0)

Config/testthat/edition 3

Encoding UTF-8

Config/roxygen2/version 8.1.0

NeedsCompilation no

Author Matthew B. Jané [aut, cre]

Maintainer Matthew B. Jané <matthewbjane@gmail.com>

Repository CRAN

Date/Publication 2026-10-21 13:50:02 UTC

Contents

as.data.frame.posc	2
autoplot.posc	3
autoplot.posc_comparison	4
posc_bins	5
posc_comp	6
posc_emp	7
posc_lm	9
posc_multi	10
posc_uni	12
predict.posc	14
print.posc	15
Index	16

as.data.frame.posc	<i>Coerce a POSC to a data frame</i>
--------------------	--------------------------------------

Description

Coerce a POSC to a data frame

Usage

```
## S3 method for class 'posc'  
as.data.frame(x, row.names = NULL, optional = FALSE, full = TRUE, ...)
```

Arguments

- x A "posc" object.
- row.names, optional Unused; included for compatibility with the generic.
- full If TRUE (default), return the curve over negative and positive differences; if FALSE, only differences ≥ 0 .
- ... Unused.

Value

A data frame with columns delta, p, lower and upper.

Examples

```
head(as.data.frame(posc_uni(r = 0.5, n = 100), full = FALSE))
```

autoplot.posc

*Plot a probability of outcome superiority curve***Description**

autoplot() returns a `ggplot2::ggplot()` object that can be modified further; plot() draws it.

Usage

```
## S3 method for class 'posc'
autoplot(
  object,
  full = FALSE,
  color = "#1F5A96",
  ci = TRUE,
  reference = FALSE,
  bins = FALSE,
  percent = FALSE,
  grid = TRUE,
  mark = NULL,
  mark_ci = FALSE,
  title = NULL,
  subtitle = NULL,
  ...
)

## S3 method for class 'posc'
plot(x, ...)
```

Arguments

object, x	A "posc" object.
full	If FALSE (default), plot differences $\Delta \geq 0$, where the curve runs from 0.5 upwards (or downwards for a negative relationship). If TRUE, plot the full sigmoid over negative and positive differences. The two halves are mirror images, since $P(-\Delta) = 1 - P(\Delta)$.
color	Colour of the curve and its confidence band.
ci	Draw the confidence band (if one is available)?
reference	For empirical curves from <code>posc_emp()</code> : also draw the bivariate-normal POSC with the same Pearson correlation, as a dashed line? Off by default.
bins	For empirical curves fitted with bins: draw the binned observed proportions as points, with whiskers showing simultaneous confidence intervals (see <code>posc_bins()</code>)? Off by default.
percent	Label the probability axis as percentages (50, 60, ...) instead of probabilities (0.5, 0.6, ...)?

grid	Draw light gridlines?
mark	Optional numeric vector of differences in the predictor to highlight. For each one, a guide line runs up from the x-axis to the curve and the probability is labelled at the point. Negative values are only shown when full = TRUE.
mark_ci	Add the confidence interval to the labels of marked points?
title, subtitle	Plot title and subtitle. None by default; supply a string, or TRUE for an automatic one (the subtitle then describes the method, correlation, sample size and band).
...	Passed from plot() to autoplot().

Value

autoplot() returns a ggplot object; plot() returns it invisibly.

Examples

```
fit <- posc_uni(r = 0.5, n = 120, sd_x = 15,
               predictor_name = "IQ", outcome_name = "job performance")
plot(fit)
plot(fit, full = TRUE)
plot(fit, mark = c(10, 20, 30))
plot(fit, mark = 15, mark_ci = TRUE, percent = TRUE)
plot(fit, title = TRUE, subtitle = TRUE)
```

```
autoplot.posc_comparison
```

Plot a comparison of POSCs

Description

Plot a comparison of POSCs

Usage

```
## S3 method for class 'posc_comparison'
autoplot(
  object,
  full = FALSE,
  colors = NULL,
  ci = TRUE,
  percent = FALSE,
  grid = TRUE,
  mark = NULL,
  mark_ci = FALSE,
  title = NULL,
  subtitle = NULL,
  ...
)
```

```
## S3 method for class 'posc_comparison'
plot(x, ...)
```

Arguments

object, x	A "posc_comparison" object from posc_comp() .
full	Plot the full sigmoid over negative and positive differences?
colors	Colours for the curves, one per curve. Defaults to a colour-blind-safe palette.
ci	Draw confidence bands (where available)?
percent	Label the probability axis as percentages instead of probabilities?
grid	Draw light gridlines?
mark, mark_ci	Differences in the predictor to highlight on every curve, and whether to add confidence intervals to their labels; see autoplot.posc() .
title, subtitle	Plot title and subtitle. None by default; supply a string, or TRUE for an automatic one (the subtitle then reports the test of equal correlations).
...	Passed from plot() to autoplot().

Value

autoplot() returns a ggplot object; plot() returns it invisibly.

Examples

```
cmp <- posc_comp(a = posc_uni(0.5, n = 100), b = posc_uni(0.3, n = 100))
plot(cmp)
plot(cmp, mark = 1, title = TRUE, subtitle = TRUE)
```

posc_bins

Binned observed proportions for an empirical POSC

Description

Returns the observed proportion of concordant pairs (the person higher on the predictor is also higher on the outcome) in bins of the absolute difference in the predictor, with simultaneous (family-wise) bootstrap confidence intervals. These are the points and whiskers drawn by `plot(fit, bins = TRUE)` for curves from [posc_emp\(\)](#); see its Details for the method.

Usage

```
posc_bins(object, level = object$level)
```

Arguments

object	A "posc" object from posc_emp() created with bins.
level	Confidence level of the simultaneous intervals. Defaults to the level used when the curve was created.

Value

A data frame with one row per bin: `delta_low` and `delta_high` (bin limits), `delta` (mean absolute difference of the pairs in the bin), `n_pairs`, `p` (observed proportion), and `lower` and `upper` (simultaneous confidence limits; NA if the curve was fitted with `n_boot = 0`). Bins that contain no pairs have NA estimates.

Examples

```
set.seed(1)
x <- rnorm(100)
y <- x + rnorm(100)
fit <- posc_emp(x, y, n_boot = 200, bins = 6)
posc_bins(fit)
```

posc_comp

Compare probability of outcome superiority curves

Description

Overlays several POSCs and tests whether their underlying correlations differ.

Usage

```
posc_comp(..., labels = NULL)
```

Arguments

`...` Two or more "posc" objects, optionally named (the names are used as labels). A single list of "posc" objects is also accepted.

`labels` Optional character vector of labels, one per curve.

Details

Under bivariate normality a POSC is determined entirely by its correlation, so curves from independent samples can be compared by testing whether the correlations are equal. For two curves this is the usual Fisher- z test,

$$z = \frac{\tanh^{-1} r_1 - \tanh^{-1} r_2}{\sqrt{1/(n_1 - 3) + 1/(n_2 - 3)}};$$

for more than two it is the homogeneity test $Q = \sum_i (n_i - 3)(z_i - \bar{z})^2$, referred to a χ^2 distribution with one fewer degrees of freedom than curves, where $\bar{z}$ is the weighted mean of the z_i .

Both tests assume the curves come from independent samples. For curves from `posc_emp()` the test compares the Pearson correlations, not the shapes of the curves. No test is reported if any curve lacks a sample size.

Value

An object of class "posc_comparison", a list with elements posc (the curves), table (a data frame of each curve's method, correlation and sample size) and test (a list with the statistic, degrees of freedom and p-value, or NULL). It has print(), plot() and `ggplot2::autoplot()` methods.

Examples

```
a <- posc_uni(r = 0.5, n = 100)
b <- posc_uni(r = 0.3, n = 120)
cmp <- posc_comp(`Structured interview` = a, `Unstructured interview` = b)
cmp
plot(cmp)
```

posc_emp

*Empirical POSC from raw data***Description**

Estimates the probability of outcome superiority curve directly from predictor and outcome scores, without assuming bivariate normality, with a bootstrap confidence band.

Usage

```
posc_emp(
  x,
  y,
  df = 3,
  level = 0.95,
  n_boot = 1000,
  max_pairs = 20000,
  max_delta = NULL,
  bins = 10,
  standardize = FALSE,
  predictor_name = "X",
  outcome_name = "Y"
)
```

Arguments

x, y	Numeric vectors of predictor and outcome scores. Pairs with a missing value in either are dropped.
df	Degrees of freedom of the spline (1 to 10).
level	Confidence level of the band.
n_boot	Number of bootstrap resamples. At least 1000 is recommended for a 95 percent band; use 0 for no band.
max_pairs	Maximum number of pairs used in each fit.

max_delta	Largest difference in the predictor to evaluate. Defaults to the 95th percentile of the absolute pairwise differences.
bins	Bins for the observed proportions: either the number of bins (default 10), which are formed to hold roughly equal numbers of pairs, or a vector of breakpoints for $ \Delta $ (in the units of the predictor, after standardizing if <code>standardize = TRUE</code>). Use 0 or NULL for no bins.
standardize	If TRUE, the predictor is standardized first so that differences are in standard-deviation units.
predictor_name, outcome_name	Labels used in printed output and plots.

Details

Pairwise data. Every pair of people (i, j) with $x_i \neq x_j$ contributes the difference $\Delta_{ij} = x_i - x_j$ and an indicator of whether the person with the higher predictor score also has the higher outcome (ties in the outcome count as 1/2). For large samples a random subset of `max_pairs` pairs is used.

Model. The curve is modelled as

$$\text{logit } P(\Delta) = \text{sign}(\Delta) g(|\Delta|),$$

where g is a natural cubic spline with `df` degrees of freedom and no intercept, whose basis is zero at $|\Delta| = 0$. This builds in the two properties every POSC must have, $P(0) = 0.5$ and $P(-\Delta) = 1 - P(\Delta)$, while leaving the shape free. The 0/1 responses are shrunk by 0.0005 towards 0.5 so that the fit exists even for near-perfect relationships; this changes the curve by a negligible amount. Increase `df` for a more flexible curve; `df = 1` makes the logit of the curve linear in Δ . Because a natural spline has zero curvature at its boundary, the curve is smooth through $\Delta = 0$. It is not forced to be monotone.

Confidence band. Pairs that share a person are not independent, so model-based standard errors from the pairwise regression would be far too small. Instead, people are resampled with replacement `n_boot` times, the curve is refitted to each resample (with the knots held fixed), and the band is formed from pointwise percentiles of the bootstrap curves. Pairs formed by a person and their own duplicate are excluded. Use `set.seed()` beforehand for reproducible results.

Binned observed proportions. To check the fitted curve against the data, the pairs with $|\Delta| \leq \text{max\_delta}$ are grouped into bins of $|\Delta|$ and the observed proportion of pairs in which the person higher on the predictor is also higher on the outcome is computed in each bin. `plot(fit, bins = TRUE)` shows these as points with whiskers, and `posc_bins()` returns them as a table. The whiskers are *simultaneous* (family-wise) confidence intervals: with probability `level` they cover all of the bin proportions at once. They are computed from the same person-level bootstrap as the curve, on the empirical-logit scale $\ell_k = \text{logit}\{(n_k \hat{p}_k + 0.5)/(n_k + 1)\}$ for a bin with n_k pairs: the interval is $\text{logit}^{-1}(\ell_k \pm c \hat{se}_k)$, where $\hat{se}_k$ is the bootstrap standard error of ℓ_k and c is the `level` quantile of $\max_k |\ell_k^* - \ell_k| / \hat{se}_k$ over the bootstrap resamples (the *max-t* or *sup-t* method). The intervals are therefore asymmetric and lie between 0 and 1.

Non-monotone relationships. A POSC averages over where in the predictor's distribution a pair of people sits. For a U-shaped relationship, the person higher on the predictor is more likely to have the higher outcome on one side of the minimum and less likely on the other, so the curve reflects the balance of the two. For a symmetric U centred on the middle of the predictor's distribution the true POSC is flat at 0.5, even though the predictor is strongly related to the outcome. The empirical curve is not constrained to be monotone and recovers such shapes; the dashed normal-theory reference, which depends only on the linear correlation, is not meaningful in that case.

Value

An object of class "posc": a list with elements

curve data frame with columns delta, p, lower, upper, over a symmetric grid of differences.

method, r, n, level, sd_x, max_delta, labels, details information about how the curve was computed.

Use [plot\(\)](#), [summary\(\)](#) and [predict\(\)](#) to work with it.

See Also

[posc_uni\(\)](#) for the curve implied by bivariate normality, which `plot(fit, reference = TRUE)` overlays as a dashed line.

Examples

```
set.seed(1)
x <- rnorm(120)
y <- exp(0.8 * x + rnorm(120, sd = 0.6)) # skewed, non-normal outcome

fit <- posc_emp(x, y, n_boot = 200) # use n_boot >= 1000 in practice
fit
plot(fit)
posc_bins(fit)
plot(fit, bins = TRUE, reference = TRUE) # add data and normal curve
```

posc_lm

POSC for the fitted values of a linear model

Description

Computes the probability of outcome superiority curve for a linear model's predictions: for two people whose predicted outcomes differ by Δ , the probability that the one predicted higher actually scores higher.

Usage

```
posc_lm(
  model,
  adjust = TRUE,
  level = 0.95,
  max_delta = NULL,
  predictor_name = NULL,
  outcome_name = NULL
)
```

Arguments

<code>model</code>	A fitted <code>stats::lm()</code> model (not <code>glm</code> or multivariate <code>mlm</code>).
<code>adjust</code>	Use the adjusted multiple correlation?
<code>level</code>	Confidence level of the band.
<code>max_delta</code>	Largest difference in the predicted outcome to evaluate. Defaults to 3 standard deviations of the fitted values.
<code>predictor_name, outcome_name</code>	Labels used in printed output and plots. <code>outcome_name</code> defaults to the name of the response variable.

Details

The fitted values are treated as a single predictor whose correlation with the observed outcome is the multiple correlation R of the model, and the curve is computed as in `posc_uni()` under bivariate normality. Differences are expressed in the units of the outcome (that is, of the fitted values). With `adjust = TRUE` (the default) R is replaced by the square root of the adjusted R^2 ; see `posc_multi()` for why.

Value

An object of class "posc": a list with elements

curve data frame with columns `delta`, `p`, `lower`, `upper`, `over` a symmetric grid of differences.

method, r, n, level, sd_x, max_delta, labels, details information about how the curve was computed.

Use `plot()`, `summary()` and `predict()` to work with it.

Examples

```
set.seed(1)
d <- data.frame(x1 = rnorm(150), x2 = rnorm(150))
d$y <- d$x1 + 0.5 * d$x2 + rnorm(150)
fit <- posc_lm(lm(y ~ x1 + x2, data = d))
fit
plot(fit)
```

posc_multi

POSC for a composite of several predictors (bivariate normal)

Description

Computes the probability of outcome superiority curve for the optimally weighted (regression) composite of several predictors, from a correlation matrix.

Usage

```
posc_multi(
  r_mat,
  n,
  outcome = ncol(r_mat),
  adjust = TRUE,
  level = 0.95,
  max_delta = 3,
  predictor_name = "Predicted Y",
  outcome_name = NULL,
  outcome_idx = NULL
)
```

Arguments

<code>r_mat</code>	Correlation matrix of the predictors and the outcome.
<code>n</code>	Sample size used to estimate <code>r_mat</code> .
<code>outcome</code>	Row/column of <code>r_mat</code> holding the outcome, as an index or a name. Defaults to the last one.
<code>adjust</code>	Use the adjusted multiple correlation? See Details.
<code>level</code>	Confidence level of the band.
<code>max_delta</code>	Largest difference in the composite (in SD units) to evaluate.
<code>predictor_name, outcome_name</code>	Labels used in printed output and plots. <code>outcome_name</code> defaults to the outcome's name in <code>r_mat</code> , if it has one.
<code>outcome_idx</code>	Deprecated; use <code>outcome</code> .

Details

The composite that best predicts the outcome correlates with it at the multiple correlation

$$R = \sqrt{\mathbf{r}_{xy}^\top \mathbf{R}_{xx}^{-1} \mathbf{r}_{xy}},$$

where $\mathbf{R}_{xx}$ is the predictor intercorrelation matrix and $\mathbf{r}_{xy}$ the predictor-outcome correlations. The POSC is then computed as in `posc_uni()`, with differences in the composite expressed in its standard-deviation units.

A sample multiple correlation is biased upwards as an estimate of the population multiple correlation, because the weights are fitted to the same sample. With `adjust = TRUE` (the default) R is replaced by the square root of the adjusted R^2 , $1 - (1 - R^2)(n - 1)/(n - k - 1)$ for k predictors, which corrects most of this bias. (How well weights estimated in this sample would predict in a new sample is lower still.)

The confidence band uses a Fisher- z interval for R with standard error $1/\sqrt{n - k - 2}$, truncated at zero, and should be treated as approximate.

Value

An object of class "posc": a list with elements

curve data frame with columns delta, p, lower, upper, over a symmetric grid of differences.

method, r, n, level, sd_x, max_delta, labels, details information about how the curve was computed.

Use [plot\(\)](#), [summary\(\)](#) and [predict\(\)](#) to work with it.

See Also

[posc_lm\(\)](#) to start from a fitted linear model instead.

Examples

```
r_mat <- matrix(c(1.0, 0.4, 0.3, 0.2,
                 0.4, 1.0, -0.3, 0.2,
                 0.3, -0.3, 1.0, 0.3,
                 0.2, 0.2, 0.3, 1.0), 4, 4,
               dimnames = list(c("X1", "X2", "X3", "Y"),
                              c("X1", "X2", "X3", "Y")))
fit <- posc_multi(r_mat, n = 100, outcome = "Y")
fit
plot(fit)
```

posc_uni

POSC from a single correlation (bivariate normal)

Description

Computes the probability of outcome superiority curve implied by a correlation between a predictor and an outcome, assuming the two are bivariate normal.

Usage

```
posc_uni(
  r,
  n = NULL,
  sd_x = 1,
  level = 0.95,
  max_delta = 3 * sd_x,
  predictor_name = "X",
  outcome_name = "Y"
)
```

Arguments

<code>r</code>	Correlation between predictor and outcome, in $(-1, 1)$.
<code>n</code>	Sample size used to estimate r . Needed for the confidence band; use NULL to draw the curve without one.
<code>sd_x</code>	Standard deviation of the predictor. The curve is expressed in the predictor's units; the default 1 gives standard-deviation units.
<code>level</code>	Confidence level of the band.
<code>max_delta</code>	Largest difference in the predictor to evaluate. Defaults to 3 standard deviations of the predictor.
<code>predictor_name, outcome_name</code>	Labels used in printed output and plots.

Details

If X and Y are bivariate normal with correlation r , then for two randomly chosen people whose predictor scores differ by Δ

$$P(Y_1 > Y_2 \mid X_1 - X_2 = \Delta) = \Phi \left(\frac{r \Delta / \sigma_X}{\sqrt{2(1 - r^2)}} \right),$$

where Φ is the standard normal distribution function and σ_X the standard deviation of the predictor. Set `sd_x` to put Δ on the raw scale of the predictor; the default `sd_x = 1` gives differences in standard-deviation units.

The curve is monotone in r , so the pointwise confidence band is obtained by evaluating it at the limits of the Fisher- z confidence interval for r .

Value

An object of class "posc": a list with elements

curve data frame with columns `delta`, `p`, `lower`, `upper`, over a symmetric grid of differences.

method, **r**, **n**, **level**, **sd_x**, **max_delta**, **labels**, **details** information about how the curve was computed.

Use [plot\(\)](#), [summary\(\)](#) and [predict\(\)](#) to work with it.

See Also

[posc_emp\(\)](#) for a curve estimated from raw data without assuming normality; [posc_comp\(\)](#) to compare curves.

Examples

```
# Differences in SD units
fit <- posc_uni(r = 0.5, n = 100)
fit
plot(fit)
```

```
# Raw units: an IQ test with SD = 15
iq <- posc_uni(r = 0.5, n = 100, sd_x = 15, predictor_name = "IQ",
              outcome_name = "job performance")
predict(iq, delta = c(5, 15, 30))
```

predict.posc

Predicted probabilities from a POSC

Description

Evaluates a probability of outcome superiority curve, with its confidence interval, at any differences in the predictor.

Usage

```
## S3 method for class 'posc'
predict(object, delta, level = object$level, ...)
```

Arguments

object	A "posc" object.
delta	Numeric vector of signed differences in the predictor (person 1 minus person 2), in the units of the predictor. Negative values give $1 - P(\Delta)$.
level	Confidence level for the interval. Defaults to the level used when the curve was created.
...	Unused.

Value

A data frame with columns delta, p (probability that person 1 has the higher outcome), lower and upper (confidence limits; NA when no interval is available).

Examples

```
fit <- posc_uni(r = 0.5, n = 100, sd_x = 15)
predict(fit, delta = c(5, 10, 15))
```

print.posc	<i>Print and summarise a POSC</i>
------------	-----------------------------------

Description

print() gives a compact description of the curve and its probabilities at a few differences in the predictor. summary() returns those probabilities as a data frame, at differences of your choosing.

Usage

```
## S3 method for class 'posc'
print(x, digits = 3, ...)

## S3 method for class 'posc'
summary(object, delta = NULL, ...)
```

Arguments

x, object	A "posc" object.
digits	Number of decimal places to print.
...	Unused.
delta	Differences in the predictor at which to report the curve. Defaults to 0.5, 1 and 2 standard deviations of the predictor.

Value

print() returns x invisibly. summary() returns a data frame as described in [predict.posc\(\)](#).

Examples

```
fit <- posc_uni(r = 0.4, n = 250)
fit
summary(fit, delta = c(0.25, 0.5, 1))
```

Index

`as.data.frame.psc`, [2](#)
`autoplot.psc`, [3](#)
`autoplot.psc()`, [5](#)
`autoplot.psc_comparison`, [4](#)

`ggplot2::autoplot()`, [7](#)
`ggplot2::ggplot()`, [3](#)

`plot()`, [9](#), [10](#), [12](#), [13](#)
`plot.psc (autoplot.psc)`, [3](#)
`plot.psc_comparison`
 (`autoplot.psc_comparison`), [4](#)
`psc_bins`, [5](#)
`psc_bins()`, [3](#), [8](#)
`psc_comp`, [6](#)
`psc_comp()`, [5](#), [13](#)
`psc_emp`, [7](#)
`psc_emp()`, [3](#), [5](#), [6](#), [13](#)
`psc_lm`, [9](#)
`psc_lm()`, [12](#)
`psc_multi`, [10](#)
`psc_multi()`, [10](#)
`psc_uni`, [12](#)
`psc_uni()`, [9–11](#)
`predict()`, [9](#), [10](#), [12](#), [13](#)
`predict.psc`, [14](#)
`predict.psc()`, [15](#)
`print.psc`, [15](#)

`set.seed()`, [8](#)
`stats::lm()`, [10](#)
`summary()`, [9](#), [10](#), [12](#), [13](#)
`summary.psc (print.psc)`, [15](#)