Package: respondeR (via r-universe)

November 11, 2026

Title Imputing Responder Proportions from Continuous Outcomes

Version 0.2.0

Description Express meta-analyses of continuous trial outcomes in terms of responder risks, following the interpretability tutorial of Thorlund, Walter, Johnston, Furukawa and Guyatt (2011) <doi:10.1002/jrsm.46>. Given the mean change, standard deviation and sample size per arm across studies, respondeR estimates the proportion of patients who cross a minimal important difference (MID) threshold under a parametric model for the change scores, and contrasts the arms as a risk difference, risk ratio, odds ratio or number needed to treat. It provides per-study (fixed- or random-effects) pooling with exact tail probabilities, a pooled-summary method anchored on the within-study mean difference, median and unweighted-mean summaries, the standardized-mean-difference to odds-ratio bridge of Anzures-Cabrera, Sarpatwari and Higgins (2011) <doi:10.1002/sim.4298>, a threshold-free common-language effect size, results across a range of thresholds, GRADE summary-of-findings rows, change-score standard deviation imputation, shared-control handling for multi-arm trials, and a point-and-click 'Shiny' application. It can also read a continuous meta-analysis already fitted with 'meta' or 'metafor' and convert it to responder risks. The framework follows the simulation study of Sofi-Mahmudi (2024) <https://hdl.handle.net/11375/30210>.

License GPL-3

URL https://github.com/choxos/respondeR, https://choxos.github.io/respondeR/

BugReports https://github.com/choxos/respondeR/issues

Encoding UTF-8

Depends R (>= 3.5)

Imports graphics, grDevices, stats, utils

Suggests shiny, shinydashboard, DT, testthat (>= 3.0.0), meta, metafor, knitr, rmarkdown

LazyData true

VignetteBuilder knitr

Config/testthat/edition 3

Config/roxygen2/version 8.1.0

NeedsCompilation no

Author Ahmad Sofi-Mahmudi [aut, cre] (ORCID: <https://orcid.org/0000-0001-6829-0823>)

Maintainer Ahmad Sofi-Mahmudi <a.sofimahmudi@gmail.com>

Repository https://adamrsophie.r-universe.dev

Date/Publication 2026-10-12 09:40:02 UTC

RemoteUrl https://github.com/cran/respondeR

RemoteRef HEAD

RemoteSha 7d7134dd517270cecd01b01964894abd0eb0e31b

Contents

as_responder_data . . . 3
estimate_change_cor . . . 4
format_responder_results . . . 4
impute_change_sd . . . 5
launch_responder_analysis . . . 6
plot.responder_curve . . . 7
responder_analysis.meta . . . 8
responder_cles . . . 12
responder_curve . . . 13
responder_from_effect . . . 15
responder_proportions . . . 17
responder_rd_individual . . . 18
responder_result . . . 20
responder_sof . . . 21
sample_responder_data . . . 22
vas_pain . . . 23

Index **25**

`as_responder_data` *Extract a responder-analysis data frame from a meta-analysis object*

Description

Pulls the per-arm summaries (`study`, `change_e`, `sd_e`, `n_e`, `change_c`, `sd_c`, `n_c`) out of a fitted continuous meta-analysis object that carries them, so the result can be passed to `responder_analysis()`. Currently supports `meta` objects of the `metacont` family (which store both arms). Studies excluded from the fitted meta-analysis (its `exclude` argument) are dropped. When several comparisons share a study label (a multi-arm trial), they get unique `study` labels and a `trial` column, so a shared control arm is split rather than double counted. Keep only the comparisons of each active arm against the control arm: `meta::pairwise()` also creates active-versus- active comparisons, which are not experimental-versus-control contrasts.

Usage

```
as_responder_data(object, ...)

## S3 method for class 'data.frame'
as_responder_data(object, ...)

## S3 method for class 'meta'
as_responder_data(object, ...)
```

Arguments

`object`	A fitted meta-analysis object, or a data frame already in the responder format.
`...`	Unused; an error if supplied.

Value

A data frame with columns `study`, `change_e`, `sd_e`, `n_e`, `change_c`, `sd_c`, `n_c` (and `trial` for multi-arm trials).

Examples

```
# A metacont-style object is just a list with these fields:
obj <- structure(
  list(
    studlab = c("A", "B"),
    n.e = c(40, 55), mean.e = c(1.1, 0.9), sd.e = c(1.2, 1.3),
    n.c = c(42, 53), mean.c = c(0.2, 0.1), sd.c = c(1.1, 1.2)
  ),
  class = c("metacont", "meta")
)
as_responder_data(obj)
```

estimate_change_cor *Estimate the baseline-to-final correlation from fully reported studies*

Description

For studies that report the baseline, final and change SDs, the implied correlation is `(sd_baseline^2 + sd_final^2 - sd_ch` (Cochrane Handbook, section 6.5.2.8). Use it (or a summary such as the median across studies) as `r` in `impute_change_sd()`.

Usage

```
estimate_change_cor(sd_baseline, sd_final, sd_change)
```

Arguments

`sd_baseline`, `sd_final`, `sd_change`	Baseline, final and change SDs.

Value

A numeric vector of correlations, truncated to `[-1, 1]`.

Examples

```
estimate_change_cor(sd_baseline = 2.1, sd_final = 2.4, sd_change = 2.2)
```

format_responder_results *Format responder-analysis results for display*

Description

Turns the numeric output of `responder_analysis()` into a compact, display-ready data frame: proportions and risk differences as percentages, the risk ratio and odds ratio with intervals, and the number needed to treat in the form of Altman (1998): `NNTB` for benefit, `NNTH` for harm, and an interval that passes through infinity when the risk-difference interval includes 0. Used by the `print()` method, the bundled Shiny app and handy for reports.

Usage

```
format_responder_results(results, digits = 1)
```

Arguments

`results`	A data frame returned by `responder_analysis()`.
`digits`	Number of decimal places (default 1).

Value

A data frame with character columns `Method`, `PE`, `PC`, `RD`, `RR`, `OR` and `NNT` (percentages for proportions and RD; ratios for RR and OR). Methods without a variance model show point estimates only.

References

Altman DG (1998). Confidence intervals for the number needed to treat. *BMJ*, 317(7168), 1309 to 1312. doi:10.1136/bmj.317.7168.1309

Examples

```
format_responder_results(responder_analysis(sample_responder_data, mid = 1))
```

`impute_change_sd` *Impute the standard deviation of a change score*

Description

Trials often report baseline and final SDs, or a confidence interval, standard error, t statistic or p-value for the mean change, rather than the SD of the change itself. These are the standard reconstructions of the Cochrane Handbook (sections 6.5.2.3 and 6.5.2.8), vectorized so they can be applied to whole columns before `responder_analysis()`:

Usage

```
impute_change_sd(
  method = c("prepost", "ci", "se", "tstat", "pvalue"),
  sd_baseline = NULL,
  sd_final = NULL,
  r = NULL,
  lower = NULL,
  upper = NULL,
  n = NULL,
  conf_level = 0.95,
  se = NULL,
  mean_change = NULL,
  t = NULL,
  p = NULL
)
```

Arguments

`method`	One of `"prepost"`, `"ci"`, `"se"`, `"tstat"`, `"pvalue"`.
`sd_baseline`, `sd_final`, `r`	Baseline and final SDs and their correlation (`"prepost"`).
`lower`, `upper`	Confidence limits for the mean change (`"ci"`).

n	Sample size ("ci", "se", "tstat", "pvalue").
conf_level	Confidence level of lower/upper (default 0.95).
se	Standard error of the mean change ("se").
mean_change	Mean change ("tstat", "pvalue").
t	One-sample t statistic of the mean change against 0 ("tstat").
p	Two-sided p-value of that t test ("pvalue").

Details

- "prepost": sqrt(sd_baseline^2 + sd_final^2 - 2 r sd_baseline sd_final), with r the correlation between baseline and final values (estimate it with estimate_change_cor() from studies that report all three SDs, and vary it in a sensitivity analysis);
- "ci": sqrt(n) (upper - lower) / (2 t_{n-1}) from a confidence interval for the mean change;
- "se": se sqrt(n);
- "tstat": |mean_change / t| sqrt(n) from a one-sample t statistic;
- "pvalue": as "tstat", with t the upper p / 2 quantile of t_{n-1} from a two-sided p-value.

Value

A numeric vector of change-score SDs.

References

Higgins JPT, Li T, Deeks JJ (editors) (2023). Chapter 6: Choosing effect measures and computing estimates of effect. In: *Cochrane Handbook for Systematic Reviews of Interventions*, version 6.4. Cochrane.

Examples

```
impute_change_sd("prepost", sd_baseline = c(2.1, 1.8), sd_final = c(2.4, 2.0), r = 0.5)
impute_change_sd("ci", lower = -1.9, upper = -0.7, n = 40)
```

---

launch_responder_analysis

Launch the Responder Analysis Shiny application

Description

Starts the bundled Shiny application, a point-and-click front end to responder_analysis(): upload data (or load the example), set the MID and direction of benefit, and view, plot and download the results.

Usage

```
launch_responder_analysis(...)
```

Arguments

...	Additional arguments passed to shiny::runApp().

Value

Called for its side effect of launching the app; invisibly returns the value of shiny::runApp().

Examples

```
launch_responder_analysis()
```

---

plot.responder_curve *Plot a responder curve*

Description

Draws the pooled risk difference against the MID threshold, one line per method, with a pointwise confidence band. Thresholds at which one study holds more than half of the weight for a pooled measure (with three or more studies) are marked with an open circle.

Usage

```
## S3 method for class 'responder_curve'
plot(x, at = NULL, ...)
```

Arguments

x	A responder_curve from responder_curve().
at	Optional MID to mark with a vertical line.
...	Passed to graphics::plot().

Value

x, invisibly.

`responder_analysis.meta` *Responder analysis of continuous trial outcomes*

Description

Converts continuous outcomes (mean change, SD and sample size per arm, across studies) into responder proportions and a range of between-arm effect measures: the risk difference (RD), risk ratio (RR), odds ratio (OR) and number needed to treat (NNT), under a parametric model for the change scores. Responders are defined by a minimal important difference (MID) threshold (the cut-point or "dichotomization" approach of Anzures-Cabrera, Sarpatwari and Higgins, 2011). For a threshold-free alternative see `responder_cles()`.

Usage

```
## S3 method for class 'meta'
responder_analysis(data, mid, ...)

## S3 method for class 'rma'
responder_analysis(
  data,
  mid = NULL,
  ...,
  baseline,
  direction = c("higher", "lower"),
  measure = NULL,
  sd_ref = NULL,
  dist = c("normal", "lognormal", "t"),
  df = NULL
)

## Default S3 method:
responder_analysis(data, ...)

responder_analysis(data, ...)

## S3 method for class 'data.frame'
responder_analysis(
  data,
  mid,
  direction = c("higher", "lower"),
  method = c("individual", "weighted", "unweighted", "median", "smd"),
  se_method = c("delta", "binomial"),
  pooling = c("fixed", "random"),
  control = c("matched", "median"),
  tau_method = c("DL", "REML"),
  dist = c("normal", "lognormal", "t"),
```

```
  df = NULL,
  mid_sd = 0,
  ci_type = c("wald", "logit"),
  ci_method = c("wald", "hksj", "hksj_adhoc"),
  conf_level = 0.95,
  rd_weights = c("sample_size", "inverse_variance"),
  ...
)
```

Arguments

data	A data frame with one row per study and columns study, change_e, sd_e, n_e, change_c, sd_c, n_c and optionally trial (see sample_responder_data); or a fitted meta-analysis object (see the "Meta-analysis objects" section).
mid	Single finite number: the minimal important difference threshold.
...	Further arguments for the method. Unknown arguments are an error.
baseline	For a metafor rma object (must be named): the control responder risk, as a probability in (0, 1) or a list with mean and sd (and optionally n). See responder_from_effect().
direction	"higher" (a larger change indicates response) or "lower".
measure	For a metafor rma object: "SMD" or "MD", overriding the measure stored in the object.
sd_ref	For a metafor rma object fitted on mean differences: the reference within-arm SD on the original scale.
dist	Change-score distribution: "normal" (default), "lognormal" or "t". Sampling variances use the distribution's own third and fourth moments.
df	Degrees of freedom when dist = "t" (a single number greater than 2).
method	Methods to compute: any of "individual", "weighted", "unweighted", "median", "smd". Defaults to the first four.
se_method	Per-study variance model for "individual": "delta" (default) propagates the sampling uncertainty of the reported mean and SD through the change-score model; "binomial" uses the pseudo-binomial p (1 - p) / n, which treats the model-based probability as an observed proportion (the version 0.1.0 default).
pooling	"fixed" (default) or "random" effects, for the "individual", "weighted" and "smd" methods.
control	Baseline-risk rule for the summary methods: "matched" (default) or "median". See the "Baseline risk" section.
tau_method	Between-study variance estimator when pooling = "random": "DL" (DerSimonian-Laird, default) or "REML" (needs the metafor package; falls back to DL with a warning if unavailable).
mid_sd	Optional standard deviation of the MID threshold; when > 0 its uncertainty is added to the effect-measure variances (once, after pooling).
ci_type	"wald" (default) or "logit". With "logit" the weighted method builds its risk-difference interval by the MOVER method from logit-scale arm intervals

(allowing for their correlation), so it stays within [-1, 1]. It applies to the weighted method and to responder_rd_individual(); pooled individual and smd intervals are formed on the scale that is pooled. A nonzero mid_sd forces Wald intervals, with a warning.

ci_method Random-effects interval method: "wald" (Normal, default), "hksj" (Hartung-Knapp-Sidik-Jonkman) or "hksj_adhoc" (HKSJ with the ad hoc floor, never narrower than Wald; recommended for few studies).

conf_level Confidence level (default 0.95).

rd_weights Fixed-effect weights for the pooled risk difference: "sample_size" (default, n_e n_c / (n_e + n_c)) or "inverse_variance".

Value

A responder_result: a data frame with one row per requested method and columns method, pooling, k, p_e, p_c, rd/rd_lb/rd_ub, rr/rr_lb/rr_ub, or/or_lb/or_ub, nnt/nnt_lb/nnt_ub, var_rd (the variance behind the risk-difference interval), and the heterogeneity statistics tau2, i2, q, q_p, pi_lb, pi_ub (prediction interval on the risk-difference scale). Proportions and risk differences are on the proportion scale; multiply by 100 for percentages. The analysis settings and the per-study statistics are kept as attributes and used by the print(), summary() and plot() methods (responder_result) and by responder_sof().

Methods

individual Dichotomize each study, then pool the per-study effect measures (fixed or random effects). The recommended method: every comparison stays within a randomized study. The per-study variance follows se_method. Ratio measures are computed from exact log probabilities, so a study deep in a probability tail cannot be pinned to an artificial ratio of 1.

weighted Pool the control-arm mean and the within-study mean difference separately (inverse variance, fixed or random effects per pooling), set the experimental mean to their sum, pool each arm's SD as the within-study pooled SD, then dichotomize. Variances come from the multivariate delta method, including the covariance between the pooled control mean and the pooled mean difference. Anchoring on the within-study difference keeps the randomized comparison intact (pooling the two arms separately, as in version 0.1.0, can produce a Simpson's-paradox bias when allocation ratios differ across studies).

unweighted, median Dichotomize the arithmetic mean or the median of the study means and SDs. Descriptive summaries with no variance model: intervals are NA.

smd Pool the standardized mean difference (Hedges' g), bridge to an odds ratio through the logistic distribution (lnOR = (pi / sqrt(3)) g; Hasselblad and Hedges, 1995; Chinn, 2000), and combine with the pooled control responder risk to recover risks. The odds ratio interval is the pooled g interval transformed (so ci_method applies); the risk difference and risk ratio intervals also carry the uncertainty of the estimated control risk. Not included by default.

Baseline risk

For the weighted method the control proportion comes, by default, from the pooled control mean (control = "matched"). Set control = "median" to take the baseline from the median control arm instead (as in the Sofi-Mahmudi 2024 simulation study); the experimental mean is then that median

control mean plus the pooled within-study mean difference. For `unweighted` and `median`, `control = "median"` takes the control proportion from the median control arm. Because the median control arm carries no sampling-variance model, `control = "median"` reports point estimates only. The `individual` and `smd` methods pool per-study contrasts and ignore `control`.

Pooling and intervals

Heterogeneity statistics (`q`, `i2`, `tau2`) are inverse-variance based and refer to the scale that was pooled: the risk difference for `individual` and Hedges' g for `smd`. With fixed effects the risk difference is pooled with sample-size weights `n_e n_c / (n_e + n_c)` by default (`rd_weights = "sample_size"`), which a single study whose responder probabilities sit near 0 or 1 (and so has a tiny model-based variance) cannot dominate; the risk ratio and odds ratio use inverse-variance weights on the log scale. A warning is raised whenever one study carries more than half of the weight for a pooled measure (with three or more studies).

For random effects, `ci_method = "hksj_adhoc"` (Hartung-Knapp-Sidik-Jonkman with the `max(1, q)` floor of Rover, Knapp and Friede, 2015) is recommended when there are few studies: plain `"hksj"` can be much narrower than the Wald interval when the studies agree closely. Prediction intervals use `t_{k-2}` and the variance that built the confidence interval, for every `tau_method`, and are `NA` for fewer than three studies.

A nonzero `mid_sd` treats the threshold as uncertain. Because the MID is shared by every study, its uncertainty is added once, after pooling, through the derivative of the pooled estimate with respect to the MID; it does not change the weights or the point estimate.

Multi-arm trials

If `data` has a `trial` column, rows with the same `trial` value and an identical control arm (`change_c`, `sd_c`, `n_c`) are treated as comparisons that share one control group, and the control sample size is divided equally among them (Cochrane Handbook, section 23.3.4) so the control patients are not counted twice. A message reports when this happens.

Meta-analysis objects

`responder_analysis()` also accepts a fitted continuous meta-analysis saved in an R object instead of a data frame. A `meta` object (a `metacont` from the **meta** package) carries the per-arm summaries, so it is converted to the data-frame form (dropping excluded studies) and analyzed with the full method menu; its pooling settings (common or random effects, a DL or REML tau-squared, a Hartung-Knapp interval with or without the ad hoc rule, and the confidence level) are used unless you override them, and settings respondeR cannot reproduce are reported in a warning. A **metafor** `rma` object carries only the pooled relative effect, so it is converted with `responder_from_effect()` and you must supply a named `baseline` control responder risk (and `sd_ref` for a mean-difference model). The object is read with base R, so the source package need not be attached.

References

Sofi-Mahmudi A (2024). Identifying an optimal strategy for converting pain as a continuous outcome to a responder analysis. Master's thesis, McMaster University. https://hdl.handle.net/11375/30210

Thorlund K, Walter SD, Johnston BC, Furukawa TA, Guyatt GH (2011). Pooling health-related quality of life outcomes in meta-analysis: a tutorial and review of methods for enhancing interpretability. *Research Synthesis Methods*, 2(3), 188 to 203. doi:10.1002/jrsm.46

Anzures-Cabrera J, Sarpatwari A, Higgins JPT (2011). Expressing findings from meta-analyses of continuous outcomes in terms of risks. *Statistics in Medicine*, 30(25), 2967 to 2985. doi:10.1002/sim.4298

Rover C, Knapp G, Friede T (2015). Hartung-Knapp-Sidik-Jonkman approach and its modification for random-effects meta-analysis with few studies. *BMC Medical Research Methodology*, 15, 99. doi:10.1186/s1287401500911

Altman DG, Deeks JJ (2002). Meta-analysis, Simpson's paradox, and the number needed to treat. *BMC Medical Research Methodology*, 2, 3. doi:10.1186/1471228823

See Also

`responder_rd_individual()`, `responder_cles()`, `responder_proportions()`, `responder_curve()`, `responder_sof()`

Examples

```
responder_analysis(sample_responder_data, mid = 1)

# Random-effects individual method with the small-k interval:
responder_analysis(sample_responder_data, mid = 1,
  method = "individual", pooling = "random", ci_method = "hksj_adhoc")
```

`responder_cles` *Common-language effect size (probabilistic index)*

Description

A threshold-free responder-type measure: the probability that a randomly chosen treated patient has a better change score than a randomly chosen control. Under a Normal model this is exact, $\mathrm{CLES} = \Phi(\delta)$ with $\delta = (\mu_e - \mu_c)/\sqrt{\sigma_e^2 + \sigma_c^2}$ (the sign is flipped when `direction = "lower"`). Per-study δ values are pooled by fixed- or random-effect inverse variance and back-transformed, so no minimal important difference is required.

Usage

```
responder_cles(
  data,
  direction = c("higher", "lower"),
  pooling = c("fixed", "random"),
  tau_method = c("DL", "REML"),
  ci_method = c("wald", "hksj", "hksj_adhoc"),
  conf_level = 0.95
)
```

Arguments

`data`	A data frame with columns `study`, `change_e`, `sd_e`, `n_e`, `change_c`, `sd_c`, `n_c`. See sample_responder_data.
`direction`	`"higher"` (a larger change is better) or `"lower"`.
`pooling`	`"fixed"` (default) or `"random"` effects.
`tau_method`	Between-study variance estimator for random effects: `"DL"` (default) or `"REML"`.
`ci_method`	Random-effects interval method: `"wald"` (default), `"hksj"` (Hartung-Knapp-Sidik-Jonkman) or `"hksj_adhoc"` (HKSJ with the ad hoc floor, never narrower than Wald; recommended for few studies).
`conf_level`	Confidence level (default `0.95`).

Value

A list with:

studies Per-study data frame: `study`, `delta`, `cles`, `cles_lb`, `cles_ub`.

cles, cles_lb, cles_ub Pooled CLES and its interval.

delta, se_delta Pooled standardized difference and the SE behind its interval.

tau2, i2, q, q_p, pi_lb, pi_ub Heterogeneity statistics; the prediction interval is back-transformed to the CLES scale.

pooling, k Settings echoed back.

References

McGraw KO, Wong SP (1992). A common language effect size statistic. *Psychological Bulletin*, 111(2), 361 to 365.

Examples

```
cles <- responder_cles(sample_responder_data)
cles$cles
```

`responder_curve` *Responder measures across a range of MID thresholds*

Description

The responder result depends on the chosen minimal important difference (MID). `responder_curve()` repeats `responder_analysis()` over a grid of thresholds and returns the pooled measures at each, so the dependence on the threshold can be reported (Thorlund and colleagues, 2011) instead of resting on one value. For the `individual` method it also records the largest share of the risk-difference weight held by one study, which flags thresholds where a single study (with responder probabilities near 0 or 1) dominates.

Usage

```
responder_curve(data, mids, method = c("individual", "weighted"), ...)
```

Arguments

data	A responder-format data frame (see responder_analysis()).
mids	Numeric vector of MID thresholds to evaluate.
method	Methods to compute at each threshold (default "individual" and "weighted").
...	Further arguments passed to responder_analysis(), such as direction, pooling or ci_method.

Details

The bands are pointwise confidence intervals at each threshold, not a simultaneous band over the whole curve.

Value

A responder_curve data frame with one row per threshold and method: mid, method, p_e, p_c, rd, rd_lb, rd_ub, rr, rr_lb, rr_ub, or, or_lb, or_ub, max_weight and max_weight_study (individual method), and warning (any warning raised at that threshold). It has a plot() method.

References

Thorlund K, Walter SD, Johnston BC, Furukawa TA, Guyatt GH (2011). Pooling health-related quality of life outcomes in meta-analysis: a tutorial and review of methods for enhancing interpretability. *Research Synthesis Methods*, 2(3), 188 to 203. doi:10.1002/jrsm.46

See Also

responder_analysis()

Examples

```
curve <- responder_curve(vas_pain, mids = seq(-2.5, -0.5, by = 0.5),
  direction = "lower")
curve[curve$method == "individual", c("mid", "rd", "rd_lb", "rd_ub", "max_weight")]
plot(curve)
```

`responder_from_effect` *Responder measures from a pooled relative effect*

Description

Converts a single pooled relative effect, a standardized mean difference (Hedges' g) or a mean difference, into absolute responder measures: the experimental responder proportion and the between-arm risk difference, risk ratio, odds ratio and number needed to treat. This is the engine used by the meta-analysis-object methods of `responder_analysis()` and is also useful on its own when you only have a pooled estimate (for example, lifted from a published meta-analysis) and a baseline control responder risk.

Usage

```
responder_from_effect(
  estimate,
  se = NULL,
  measure = c("SMD", "MD"),
  baseline,
  mid = NULL,
  direction = c("higher", "lower"),
  sd_ref = NULL,
  dist = c("normal", "lognormal", "t"),
  df = NULL,
  ci = NULL,
  conf_level = 0.95,
  label = NULL
)
```

Arguments

`estimate`	Single number: the pooled relative effect (Hedges' g for `measure = "SMD"`, or a mean difference for `measure = "MD"`), experimental minus control.
`se`	Standard error of `estimate`. If `NULL` and `ci` is also `NULL`, only point estimates are returned (no intervals).
`measure`	`"SMD"` (default) or `"MD"`.
`baseline`	The control (baseline) responder risk, supplied as either a single probability in `(0, 1)`, or a list with `mean` and `sd` (and optionally `n`) giving the control-arm mean change and SD, from which the control responder proportion is computed at the threshold.
`mid`	Minimal important difference threshold. Required when `baseline` is given as a mean and SD.
`direction`	`"higher"` (a larger change indicates response) or `"lower"`.
`sd_ref`	Reference within-arm standard deviation on the original scale. Required for `measure = "MD"`.

dist, df	Change-score distribution and degrees of freedom for the control responder probability when baseline is a mean and SD. The MD conversion requires dist = "normal".
ci	Optional confidence interval for estimate (a length-2 numeric vector), for example a t-based interval from a fitted model. When given, it sets the critical value, so the interval method of the source analysis is kept.
conf_level	Confidence level (default 0.95); used when ci is not supplied.
label	Optional label placed in the method column of the result.

Details

estimate is the effect as experimental minus control on the outcome's own scale, as reported by **meta** and **metafor**; direction says whether a higher change means response, and the effect is oriented accordingly (so a pain reduction, a negative difference with direction = "lower", raises the responder risk).

For a standardized mean difference the effect is mapped to an odds ratio through the logistic distribution, lnOR = (pi / sqrt(3)) * g (Hasselblad and Hedges, 1995; Chinn, 2000; Anzures-Cabrera, Sarpatwari and Higgins, 2011), and combined with the control proportion to recover risks. For a mean difference the experimental proportion is obtained by shifting the control responder probability along the Normal cut-point (Thorlund and colleagues, 2011), which requires a reference within-arm standard deviation sd_ref; this conversion assumes Normal change scores.

When baseline is a fixed probability the intervals are conditional on that assumed control risk (as in a GRADE summary-of-findings table). When it is estimated (a mean and SD with n), its sampling uncertainty is propagated jointly with that of the effect.

Value

A one-row responder_result in the same schema as responder_analysis().

References

Thorlund K, Walter SD, Johnston BC, Furukawa TA, Guyatt GH (2011). Pooling health-related quality of life outcomes in meta-analysis: a tutorial and review of methods for enhancing interpretability. *Research Synthesis Methods*, 2(3), 188 to 203. doi:10.1002/jrsm.46

Anzures-Cabrera J, Sarpatwari A, Higgins JPT (2011). Expressing findings from meta-analyses of continuous outcomes in terms of risks. *Statistics in Medicine*, 30(25), 2967 to 2985. doi:10.1002/sim.4298

Chinn S (2000). A simple method for converting an odds ratio to effect size for use in meta-analysis. *Statistics in Medicine*, 19(22), 3127 to 3131.

See Also

responder_analysis(), responder_sof()

Examples

```
# A pooled SMD of 0.5 (SE 0.1) against a 25% control responder rate:
responder_from_effect(0.5, se = 0.1, measure = "SMD", baseline = 0.25)

# A pooled mean difference of 1.2 (SE 0.3), control mean 0.2, SD 1.3, MID 1:
responder_from_effect(1.2, se = 0.3, measure = "MD",
  baseline = list(mean = 0.2, sd = 1.3), mid = 1, sd_ref = 1.3)

# Pain (lower is better): a negative SMD is a benefit.
responder_from_effect(-0.5, se = 0.1, baseline = 0.25, direction = "lower")
```

---

responder_proportions *Responder proportions from continuous arm summaries*

Description

Estimates, for each study arm, the probability that a patient's change score crosses the minimal important difference (MID) threshold under a parametric model for the change scores, together with a delta-method (sampling) variance for that probability.

Usage

```
responder_proportions(
  change,
  sd,
  n,
  mid,
  direction = c("higher", "lower"),
  dist = c("normal", "lognormal", "t"),
  df = NULL
)
```

Arguments

change	Numeric vector of mean change scores.
sd	Numeric vector of standard deviations (> 0).
n	Numeric vector of sample sizes (>= 2).
mid	Single finite number: the minimal important difference threshold.
direction	"higher" (a larger change indicates response) or "lower".
dist	Change-score distribution: "normal" (default), "lognormal" or "t".
df	Degrees of freedom when dist = "t" (a single number greater than 2).

Details

The probability is computed from the requested tail on the log scale, so very small probabilities keep full relative precision. The variance propagates the sampling uncertainty of the arm mean and SD, using the third and fourth moments of the chosen distribution (for the Normal model this is phi(a)^2 (1/n + a^2 / (2 (n - 1))) with a = (mu - mid) / sigma).

Value

A data frame with one row per input element and columns `p` (responder probability) and `var_p` (delta-method variance).

References

Thorlund K, Walter SD, Johnston BC, Furukawa TA, Guyatt GH (2011). Pooling health-related quality of life outcomes in meta-analysis: a tutorial and review of methods for enhancing interpretability. *Research Synthesis Methods*, 2(3), 188 to 203. doi:10.1002/jrsm.46

Anzures-Cabrera J, Sarpatwari A, Higgins JPT (2011). Expressing findings from meta-analyses of continuous outcomes in terms of risks. *Statistics in Medicine*, 30(25), 2967 to 2985. doi:10.1002/sim.4298

Examples

```
responder_proportions(
  change = c(0.96, 0.79, 1.02), sd = c(1.26, 1.28, 1.34),
  n = c(43, 139, 156), mid = 1
)
```

`responder_rd_individual`

Per-study responder risk differences

Description

Dichotomizes each study at the MID threshold and returns the per-study responder risk difference (experimental minus control) with a confidence interval. Building block for the `"individual"` method of `responder_analysis()`; also feeds the forest plot and per-study table.

Usage

```
responder_rd_individual(
  data,
  mid,
  direction = c("higher", "lower"),
  se_method = c("delta", "binomial"),
  conf_level = 0.95,
  dist = c("normal", "lognormal", "t"),
  df = NULL,
  mid_sd = 0,
  ci_type = c("wald", "logit")
)
```

Arguments

data	A data frame in the responder format with one row per study (columns study, change_e, sd_e, n_e, change_c, sd_c, n_c, and optionally trial); see sample_responder_data.
mid	Single finite number: the minimal important difference threshold.
direction	"higher" (a larger change indicates response) or "lower".
se_method	Per-study variance model for "individual": "delta" (default) propagates the sampling uncertainty of the reported mean and SD through the change-score model; "binomial" uses the pseudo-binomial p (1 - p) / n, which treats the model-based probability as an observed proportion (the version 0.1.0 default).
conf_level	Confidence level (default 0.95).
dist	Change-score distribution: "normal" (default), "lognormal" or "t". Sampling variances use the distribution's own third and fourth moments.
df	Degrees of freedom when dist = "t" (a single number greater than 2).
mid_sd	Optional standard deviation of the MID threshold; when > 0 its uncertainty is added to the effect-measure variances (once, after pooling).
ci_type	"wald" (default) or "logit". With "logit" the weighted method builds its risk-difference interval by the MOVER method from logit-scale arm intervals (allowing for their correlation), so it stays within [-1, 1]. It applies to the weighted method and to responder_rd_individual(); pooled individual and smd intervals are formed on the scale that is pooled. A nonzero mid_sd forces Wald intervals, with a warning.

Details

With ci_type = "logit" each interval is built by Newcombe's MOVER method from logit-scale arm intervals and stays within [-1, 1]; a nonzero mid_sd adds the threshold uncertainty (which correlates the two arms) on the Wald scale instead.

Value

A data frame with one row per study and columns study, p_e, p_c, rd, se, ci_lb, ci_ub (proportion scale).

See Also

responder_analysis()

Examples

```
responder_rd_individual(sample_responder_data, mid = 1)
```

`responder_result` *Methods for responder analysis results*

Description

`responder_analysis()` and `responder_from_effect()` return a `responder_result`: a data frame with one row per method (so `$`, `[` and column access work as usual) that also keeps the analysis settings and the per-study statistics. `print()` shows a formatted table with the number needed to treat; `summary()` adds the settings, heterogeneity statistics and per-study weights; `plot()` draws a forest plot of the per-study and pooled risk differences; `as.data.frame()` returns the plain numeric table.

Usage

```
## S3 method for class 'responder_result'
print(x, digits = 1, ...)

## S3 method for class 'responder_result'
x[i, j, drop = FALSE]

## S3 method for class 'responder_result'
summary(object, ...)

## S3 method for class 'summary.responder_result'
print(x, digits = 1, ...)

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

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

Arguments

`x`, `object`	A `responder_result`.
`digits`	Decimal places for percentages.
`...`	Unused (for `plot()`, passed to `graphics::plot()`).
`i`, `j`	Row and column indices.
`drop`	Passed to the data-frame method.
`row.names`, `optional`	Passed to `as.data.frame()`.

Value

`print()` and `plot()` return `x` invisibly; `summary()` returns a `summary.responder_result` list; `as.data.frame()` returns a data frame.

Examples

```
res <- responder_analysis(sample_responder_data, mid = 1)
res
summary(res)
plot(res)
head(as.data.frame(res), 2)
```

`responder_sof` *GRADE summary-of-findings rows from a responder analysis*

Description

Expresses pooled responder results as absolute effects per 1000 patients at an assumed control (baseline) responder risk, as in a GRADE summary-of-findings table: the control risk, the corresponding intervention risk and the difference, each with a confidence interval, and the relative effect. As GRADE recommends, the absolute effect is derived from the pooled relative effect applied to a representative control risk (Guyatt and colleagues, 2013).

Usage

```
responder_sof(
  x,
  control_risk = NULL,
  measure = c("OR", "RR", "RD"),
  per = 1000
)
```

Arguments

`x`	A result from `responder_analysis()` or `responder_from_effect()`.
`control_risk`	Assumed control responder risk(s) in `(0, 1)`. Defaults to the median control-arm responder probability stored in `x`.
`measure`	Relative measure used for the conversion: `"OR"` (default), `"RR"` or `"RD"`.
`per`	Denominator for the absolute effects (default 1000).

Details

The default relative measure is the odds ratio. The cut-point model behind a responder analysis implies a roughly constant effect on the logit (probit) scale, not a constant risk ratio, so the odds ratio transfers across baseline risks best and always yields a valid risk. When responders are common in the control arms, a pooled responder risk ratio is compressed toward 1 and understates the absolute effect at lower baseline risks; use `measure = "RR"` when a risk ratio is required.

The assumed control risk defaults to the median control-arm responder probability across the studies at the analysis MID. Supply one or more values in `control_risk` (for example low, moderate and high baseline risk) to get a row for each. The intervention risk is `control_risk * RR` (`measure = "RR"`), `plogis(qlogis(control_risk) + log(OR))` (`"OR"`) or `control_risk + RD` (`"RD"`); its

interval applies the measure's bounds to the same control risk. A risk-ratio conversion that would exceed a risk of 1 is reported as NA with a warning (use the odds ratio instead). Certainty ratings are a reviewer judgment and are not produced here.

Value

A `responder_sof` data frame with one row per method and control risk: `method`, `measure`, `control_risk`, the per-per columns `control`, `intervention`, `intervention_lb`, `intervention_ub`, `difference`, `difference_lb`, `difference_ub`, and the relative effect `relative`, `relative_lb`, `relative_ub`. Its `print()` method shows the GRADE wording ("255 more per 1000 (187 more to 324 more)").

References

Guyatt GH, Oxman AD, Santesso N, and colleagues (2013). GRADE guidelines: 12. Preparing summary of findings tables: binary outcomes. *Journal of Clinical Epidemiology*, 66(2), 158 to 172. doi:10.1016/j.jclinepi.2012.01.012

See Also

`responder_analysis()`

Examples

```
res <- responder_analysis(sample_responder_data, mid = 1, method = "individual")
responder_sof(res)
responder_sof(res, control_risk = c(0.1, 0.3))
responder_sof(res, measure = "RR")
```

---

sample_responder_data *Example responder-analysis dataset*

Description

A small illustrative dataset of three trials reporting continuous change scores per arm. Used in examples, the bundled Shiny app and the package tests. Values are fictional but plausible.

Usage

```
sample_responder_data
```

Format

A data frame with 3 rows and 7 columns:

study Study identifier.

change_e Mean change in the experimental arm.

sd_e Standard deviation of change in the experimental arm.

n_e Sample size of the experimental arm.

change_c Mean change in the control arm.

sd_c Standard deviation of change in the control arm.

n_c Sample size of the control arm.

Examples

```
responder_analysis(sample_responder_data, mid = 1)
```

`vas_pain` *VAS pain change scores from a spinal-health exercise meta-analysis*

Description

Per-comparison change in pain on a 0-10 cm visual analogue scale (VAS) for an exercise-therapy arm versus a control arm: the 20 comparisons, from 18 randomized trials, pooled for the VAS outcome by Li, Bao, Wang and Zhao (2025). The change scores are post minus baseline VAS, so a more negative value is a larger pain reduction; analyze with `direction = "lower"` and a negative MID equal to the required reduction (for example `mid = -1.5` for a 1.5 cm responder threshold).

Usage

```
vas_pain
```

Format

A data frame with 20 rows and 8 columns:

study Comparison label (first author and year).

change_e Mean VAS change in the exercise (experimental) arm.

sd_e Standard deviation of the change in the exercise arm.

n_e Exercise-arm sample size.

change_c Mean VAS change in the control arm.

sd_c Standard deviation of the change in the control arm.

n_c Control-arm sample size.

trial Trial label; equal for the comparisons of one multi-arm trial.

Details

Two trials had three arms and contribute two exercise-versus-control comparisons each, reusing one control group: Farzaneh 2022 (rows "Farzaneh a" and "Farzaneh b") and Peng Qiu 2014 ("Peng Qiu a" and "Peng Qiu b"). The `trial` column marks them, and respondeR divides each shared control arm among its comparisons (Cochrane Handbook, section 23.3.4) so those patients are not counted twice. The values are as extracted by the source meta-analysis; its figure lists 14 patients in both Farzaneh exercise arms, whereas the trial report gives 12 for one of them. The published extraction is kept so the dataset reproduces the source forest plot.

Source

Li Z, Bao Z, Wang S, Zhao M (2025). Meta-analysis of the best exercise mode and dose study for improving spinal health. *Frontiers in Sports and Active Living*, 7, 1614906. doi:10.3389/fspor.2025.1614906. Figure 3 (VAS pain). Reproduced under the Creative Commons Attribution License (CC BY 4.0); the original authors and journal are credited as required.

Examples

```
# Proportion achieving at least a 1.5 cm VAS reduction (responder),
# exercise versus control, random effects with the small-k HKSJ interval.
responder_analysis(vas_pain, mid = -1.5, direction = "lower",
  pooling = "random", ci_method = "hksj_adhoc")
```

Index

∗ **datasets**
    sample_responder_data, 22
    vas_pain, 23
[.responder_result *(*responder_result*)*, 20

as.data.frame(), *20*
as.data.frame.responder_result *(*responder_result*)*, 20
as_responder_data, 3

estimate_change_cor, 4
estimate_change_cor(), *6*

format_responder_results, 4

graphics::plot(), *7*, *20*

impute_change_sd, 5
impute_change_sd(), *4*

launch_responder_analysis, 6

plot.responder_curve, 7
plot.responder_result *(*responder_result*)*, 20
print.responder_result *(*responder_result*)*, 20
print.summary.responder_result *(*responder_result*)*, 20

responder_analysis *(*responder_analysis.meta*)*, 8
responder_analysis(), *3–6*, *13–16*, *18–22*
responder_analysis.meta, 8
responder_cles, 12
responder_cles(), *8*, *12*
responder_curve, 13
responder_curve(), *7*, *12*
responder_from_effect, 15
responder_from_effect(), *9*, *11*, *20*, *21*
responder_proportions, 17
responder_proportions(), *12*
responder_rd_individual, 18
responder_rd_individual(), *10*, *12*, *19*
responder_result, *10*, 20
responder_sof, 21
responder_sof(), *10*, *12*, *16*

sample_responder_data, *9*, *13*, *19*, 22
shiny::runApp(), *7*
summary.responder_result *(*responder_result*)*, 20

vas_pain, 23