

Package: baymedr (via r-universe)

October 10, 2026

Title Computation of Bayes Factors for Common Biomedical Designs

Version 0.2

Date 2026-05-04

Description BAYesian inference for MEDical designs in R. Functions for the computation of Bayes factors for common biomedical research designs. Implemented are functions to test the equivalence (equiv_bf), non-inferiority (infer_bf), and superiority (super_bf) of an experimental group compared to a control group on a continuous outcome measure, as well as functions for simulating survival data and calculating a Bayes factor for Cox proportional hazards models. Bayes factors for these tests can be computed based on raw data or summary statistics.

Depends R (>= 3.4.0)

Imports doParallel, foreach, methods, parallel, pso, rms, stats, stringr, survival

Suggests knitr, rmarkdown, testthat

License GPL-3

Encoding UTF-8

Language en-US

URL <https://github.com/maxlinde/baymedr>

BugReports <https://github.com/maxlinde/baymedr/issues>

Config/roxygen2/version 8.0.0

Config/pak/sysreqs cmake make libicu-dev libuv1-dev

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

Date/Publication 2026-05-04 14:22:41 UTC

RemoteUrl <https://github.com/maxlinde/baymedr>

RemoteRef HEAD

RemoteSha 86f896485e56c868678a87b3ac5f85f1fb2dcc99

Contents

baymedr-package	2
coxph_bf	3
coxph_data_sim	5
equiv_bf	7
get_bf	11
infer_bf	12
model-classes	17
super_bf	17
Index	22

baymedr-package	<i>baymedr: Computation of Bayes Factors for Common Biomedical Designs</i>
-----------------	--

Description

baymedr provides functions for the computation of Bayes factors for common biomedical research designs.

Details

Package: baymedr
 Type: Package
 Version: 0.2
 License: GPL-3
 Date: 2022-09-01

Author(s)

Maintainer: Maximilian Linde – <maximilian.linde.92@gmail.com>

Authors:

- Maximilian Linde (aut, cre)
- Don van Ravenzwaaij (aut)
- Jorge N. Tendeiro (aut)
- Quentin F. Gronau (ctb)

See Also

Useful links:

- <https://github.com/maxlinde/baymedr>
- Report bugs at <https://github.com/maxlinde/baymedr/issues>

coxph_bf

*Bayes factor for Cox proportional hazards regression***Description**

`coxph_bf` computes a Bayes factor for Cox proportional hazards regression models with one dichotomous independent variable.

Usage

```
coxph_bf(
  data,
  null_value = 0,
  alternative = "two.sided",
  direction = NULL,
  prior_mean = 0,
  prior_sd = 1
)
```

Arguments

<code>data</code>	A data.frame or a list resulting from calling <code>coxph_data_sim</code> . In the former case, the first column represents the survival/censoring times, the second column indicates whether an event (e.g., death) has happened (1) or not (0), and the third column represents the dummy-coded independent variable, where the control group is coded with 0 and the experimental group with 1.
<code>null_value</code>	The value of the point null hypothesis for the beta coefficient. The default is a null value of 0.
<code>alternative</code>	A string specifying whether the alternative hypothesis is two-sided ("two.sided"; the default) or one-sided ("one.sided").
<code>direction</code>	A string specifying the direction of the one-sided alternative hypothesis. This is ignored if <code>alternative = "two.sided"</code> . Possible options are "low" and "high".
<code>prior_mean</code>	Mean of the Normal prior for the beta parameter. The default is a mean of 0.
<code>prior_sd</code>	Standard deviation of the Normal prior for the beta parameter. The default is a standard deviation of 1.

Details

The Cox proportional hazards model has the following hypotheses: The null hypothesis (i.e., H_0) states that the population hazard ratio between the experimental (e.g., a new medication) and the control group (e.g., a placebo or an already existing medication) is equal to 1 (i.e., $\beta = 0$). The alternative hypothesis can be two-sided or one-sided (either negative or positive).

Since the main goal of `coxph_bf` is to establish that the hazard ratio is not equal to 1, the resulting Bayes factor quantifies evidence in favor of the alternative hypothesis. For a two-sided alternative hypothesis, we have BF_{10} ; for a negative one-sided alternative hypothesis, we have BF_{-0} ; and for

a positive one-sided alternative hypothesis, we have $BF+0$. Evidence for the null hypothesis can easily be calculated by taking the reciprocal of the original Bayes factor (i.e., $BF_{01} = 1 / BF_{10}$). Quantification of evidence in favor of the null hypothesis is logically sound and legitimate within the Bayesian framework (see e.g., van Ravenzwaaij et al., 2019).

For the calculation of the Bayes factor, a Normal prior density is chosen for beta under the alternative hypothesis. The arguments `prior_mean` and `prior_sd` specify the mean and standard deviation of the Normal prior, respectively. By adjusting the Normal prior, different ranges of expected effect sizes can be emphasized. The default is a Normal prior with a mean of 0 and a standard deviation of 1.

Note that at the moment the model specifications are limited. That is, it is only possible to have a single dichotomous independent variable. Further, at the moment only a Normal prior is supported. Lastly, only the Efron partial likelihood and not the many other options are supported.

`coxph_bf` creates an S4 object of class `baymedrCoxProportionalHazards`, which has multiple slots/entries (e.g., `prior`, Bayes factor, etc.; see `Value`). If it is desired to store or extract solely the Bayes factor, the user can do this with `get_bf`, by setting the S4 object as an argument (see `Examples`).

Value

An S4 object of class `baymedrCoxProportionalHazards` is returned. Contained are a description of the model and the resulting Bayes factor:

- `test`: The type of analysis
- `hypotheses`: A statement of the hypotheses
 - `h0`: The null hypothesis
 - `h1`: The alternative hypothesis
- `prior`: The parameters of the Normal prior on beta
- `bf`: The resulting Bayes factor

A summary of the model is shown by printing the object.

References

Harrell, F. R. (2015). Regression modeling strategies: With applications to linear models, logistic regression, and survival analysis (2nd ed.). Springer.

van Ravenzwaaij, D., Monden, R., Tendeiro, J. N., & Ioannidis, J. P. A. (2019). Bayes factors for superiority, non-inferiority, and equivalence designs. *BMC Medical Research Methodology*, 19, 71.

See Also

[coxph_data_sim](#).

Examples

```
# Load aml dataset from the survival R package.
data <- survival::aml
data$x <- ifelse(test = data$x == "Maintained",
               yes = 0,
               no = 1)
```

```

names(data) <- c("time", "event", "group")

# Assign model to variable.
coxph_mod <- coxph_bf(data = data,
                      null_value = 0,
                      alternative = "one.sided",
                      direction = "high",
                      prior_mean = 0,
                      prior_sd = 1)

# Extract Bayes factor from variable.
get_bf(coxph_mod)

```

coxph_data_sim	<i>Simulate data for Cox proportional hazards regression</i>
----------------	--

Description

[coxph_data_sim](#) simulates data for Cox proportional hazards regression models with one dichotomous independent variable based on summary statistics.

Usage

```

coxph_data_sim(
  n_data = 1,
  ns_c,
  ns_e,
  ne_c,
  ne_e,
  cox_hr,
  cox_hr_ci_level = 0.95,
  max_t = 100,
  cores = 1,
  ...
)

```

Arguments

<code>n_data</code>	The number of datasets to be simulated. The default is 1.
<code>ns_c</code>	Sample size of the control condition.
<code>ns_e</code>	Sample size of the experimental condition.
<code>ne_c</code>	Number of events (e.g., death) in the control condition.
<code>ne_e</code>	Number of events (e.g., death) in the experimental condition.
<code>cox_hr</code>	A numeric vector of length 3, indicating the hazard ratio between the experimental and control conditions based on a Cox proportional hazards regression model, the lower boundary of the x of the hazard ratio between the experimental

	and control conditions based on a Cox proportional hazards regression model, and the upper boundary of the x control conditions based on a Cox proportional hazards regression model, respectively. The hazard ratio must be provided. The confidence interval boundaries are optional; if missing they should be given as NA.
cox_hr_ci_level	Confidence level of the x hazard ratio between the experimental and control conditions based on a Cox proportional hazards regression model. The default is 0.95.
max_t	The maximum allowed survival/censoring time. The default is 100.
cores	The number of cores to be used in the data simulation process. The default is 1. Note that it is only useful to use more than 1 core if more than 1 dataset is simulated; ideally, n_data should be a multiple of cores.
...	Arguments passed to the control argument of <code>psoptim</code> (e.g. <code>maxit</code> , <code>maxit.stagnate</code>). Be aware that <code>coxph_data_sim</code> uses default values that are not the default in <code>psoptim</code> . Specifically, <code>coxph_data_sim</code> uses <code>maxit = 5000</code> , and <code>maxit.stagnate = ceiling(maxit / 5)</code> .

Details

Particle swarm optimization, as implemented by `psoptim` is used to simulate one or multiple datasets that match certain summary statistics. The algorithm uses as many parameters as there cases in the dataset that is to be simulated. Therefore, using `coxph_data_sim` becomes more time-consuming the larger the sample size.

The relevant summary statistics that are used in the optimization process are:

- `cox_hr`
 - Hazard ratio between the experimental and control conditions based on a Cox proportional hazards regression model.
 - Lower boundary of the x of the hazard ratio between the experimental and control conditions based on a Cox proportional hazards regression model.
 - Upper boundary of the x confidence interval of the hazard ratio between the experimental and control conditions based on a Cox proportional hazards regression model.

`coxph_data_sim` creates a list with as many elements as specified by the argument `n_data`. Each element consists of a list that entails the resulting simulated data and the optimization results of the data simulation process.

Value

A list of length `n_data` is returned. Each element of that list contains one simulated dataset and information about the optimization process:

- `data`: A `data.frame` containing the following columns:
 - `time`: Survival/censoring times.
 - `event`: Indication of whether an event happened (1) or not (0).
 - `group`: Indication of whether case belongs to control condition (0) or experimental condition (1).

- `optim`: Results of particle swarm optimization. See the Value section in [psoptim](#)

References

Harrell, F. R. (2015). Regression modeling strategies: With applications to linear models, logistic regression, and survival analysis (2nd ed.). Springer.

Kennedy, J., & Eberhart, R. (1995). Particle swarm optimization. *Proceedings of ICNN'95 - International Conference on Neural Networks*, 4, 1942-1948.

Shi, Y., & Eberhart, R. (1998). A modified particle swarm optimizer. *1998 IEEE International Conference on Evolutionary Computation Proceedings. IEEE World Congress on Computational Intelligence*, 69-73.

See Also

[coxph_bf](#) and [psoptim](#).

Examples

```
# Pretend we extracted the following summary statistics from an article.
ns_c <- 20
ns_e <- 56
ne_c <- 18
ne_e <- 40
cox_hr <- c(0.433, 0.242, 0.774)
cox_hr_ci_level <- 0.95

# We want to simulate 3 datasets. We do not need a very precise match of the
# summary statistics to the real summary statistics. Therefore, for
# demonstration purposes we only use 1/200 of the default number of
# optimization iterations (i.e., (1 / 200) * 5000).
sim_data <- coxph_data_sim(n_data = 3,
                           ns_c = ns_c,
                           ns_e = ns_e,
                           ne_c = ne_c,
                           ne_e = ne_e,
                           cox_hr = cox_hr,
                           cox_hr_ci_level = cox_hr_ci_level,
                           max_t = 100,
                           maxit = 25)
```

equiv_bf

Bayes factor for equivalence designs

Description

[equiv_bf](#) computes a Bayes factor for equivalence designs with a continuous dependent variable.

Usage

```
equiv_bf(
  x = NULL,
  y = NULL,
  n_x = NULL,
  n_y = NULL,
  mean_x = NULL,
  mean_y = NULL,
  sd_x = NULL,
  sd_y = NULL,
  ci_margin = NULL,
  ci_level = NULL,
  interval = 0,
  interval_std = TRUE,
  prior_scale = 1/sqrt(2)
)
```

Arguments

x	A numeric vector of observations for the control group.
y	A numeric vector of observations for the experimental group.
n_x	A numeric vector of length one, specifying the sample size of the control group.
n_y	A numeric vector of length one, specifying the sample size of the experimental group.
mean_x	A numeric vector of length one, specifying the mean of the dependent variable in the control group.
mean_y	A numeric vector of length one, specifying the mean of the dependent variable in the experimental group.
sd_x	A numeric vector of length one, specifying the standard deviation of the dependent variable in the control group. Only sd_x and sd_y OR ci_margin and ci_level should be defined (see Details).
sd_y	A numeric vector of length one, specifying the standard deviation of the dependent variable in the experimental group. Only sd_x and sd_y OR ci_margin and ci_level should be defined (see Details).
ci_margin	A numeric vector of length one, specifying the margin of the confidence interval (i.e., the width of the confidence interval divided by 2) of the mean difference on the dependent variable between the experimental and control groups. The value should be a positive number Only sd_x and sd_y OR ci_margin and ci_level should be defined (see Details).
ci_level	A numeric vector of length one, specifying the confidence level of ci_margin. The value must be between 0 and 1 (e.g., 0.95 for a 95% confidence interval). Only sd_x and sd_y OR ci_margin and ci_level should be defined (see Details).
interval	A numeric vector of length one or two, specifying the boundaries of the equivalence interval. If a numeric vector of length one is specified, a symmetric

	equivalence interval will be used (e.g., a 0.1 is equivalent to $c(-0.1, 0.1)$). A numeric vector of length two provides the possibility to specify an asymmetric equivalence interval (e.g., $c(-0.1, 0.2)$). The default is 0, indicating a point null hypothesis rather than an interval (see Details).
interval_std	A logical vector of length one, specifying whether the equivalence interval (i.e., interval) is given in standardized (TRUE; the default) or unstandardized (FALSE) units.
prior_scale	A numeric vector of length one, specifying the scale of the Cauchy prior distribution for the effect size under the alternative hypothesis (see Details). The default value is $r = 1 / \sqrt{2}$.

Details

The equivalence design has the following hypotheses: The null hypothesis (i.e., H_0) states that the population means of the experimental group (e.g., a new medication) and the control group (e.g., a placebo or an already existing medication) are (practically) equivalent; the alternative hypothesis (i.e., H_1) states that the population means of the two groups are not equivalent. The dependent variable must be continuous.

Since the main goal of `equiv_bf` is to establish equivalence, the resulting Bayes factor quantifies evidence in favor of the null hypothesis (i.e., BF_{01}). Evidence for the alternative hypothesis can easily be calculated by taking the reciprocal of the original Bayes factor (i.e., $BF_{10} = 1 / BF_{01}$). Quantification of evidence in favor of the null hypothesis is logically sound and legitimate within the Bayesian framework (see e.g., van Ravenzwaaij et al., 2019).

`equiv_bf` can be utilized to calculate a Bayes factor based on raw data (i.e., if arguments `x` and `y` are defined) or summary statistics (i.e., if arguments `n_x`, `n_y`, `mean_x`, and `mean_y` are defined). In the latter case, either values for the arguments `sd_x` and `sd_y` **OR** `ci_margin` and `ci_level` can be supplied. Arguments with 'x' as a name or suffix correspond to the control group, whereas arguments with 'y' as a name or suffix correspond to the experimental group.

The equivalence interval can be specified with the argument `interval`. However, it is not compulsory to specify an equivalence interval (see van Ravenzwaaij et al., 2019). The default value of the argument `interval` is 0, indicating a point null hypothesis. If an interval is preferred, the argument `interval` can be set in two ways: A *symmetric* interval can be defined by either specifying a numeric vector of length one (e.g., 0.1, which is converted to $c(-0.1, 0.1)$) or a numeric vector of length two (e.g., $c(-0.1, 0.1)$); an *asymmetric* interval can be defined by specifying a numeric vector of length two (e.g., $c(-0.1, 0.2)$). It can be specified whether the equivalence interval (i.e., interval) is given in standardized or unstandardized units with the `interval_std` argument, where TRUE, corresponding to standardized units, is the default.

For the calculation of the Bayes factor, a Cauchy prior density centered on 0 is chosen for the effect size under the alternative hypothesis. The standard Cauchy distribution, with a location parameter of 0 and a scale parameter of 1, resembles a standard Normal distribution, except that the Cauchy distribution has less mass at the center but heavier tails (Liang et al., 2008; Rouder et al., 2009). The argument `prior_scale` specifies the width of the Cauchy prior, which corresponds to half of the interquartile range. Thus, by adjusting the Cauchy prior scale with `prior_scale`, different ranges of expected effect sizes can be emphasized. The default prior scale is set to $r = 1 / \sqrt{2}$.

`equiv_bf` creates an S4 object of class `baymedrEquivalence`, which has multiple slots/entries (e.g., type of data, prior scale, Bayes factor, etc.; see Value). If it is desired to store or extract solely

the Bayes factor, the user can do this with `get_bf`, by setting the S4 object as an argument (see Examples).

Value

An S4 object of class `baymedrEquivalence` is returned. Contained are a description of the model and the resulting Bayes factor:

- test: The type of analysis
- hypotheses: A statement of the hypotheses
 - h0: The null hypothesis
 - h1: The alternative hypothesis
- interval: Specification of the equivalence interval in standardized and unstandardized units
 - lower_std: The standardized lower boundary of the equivalence interval
 - upper_std: The standardized upper boundary of the equivalence interval
 - lower_unstd: The unstandardized lower boundary of the equivalence interval
 - upper_unstd: The unstandardized upper boundary of the equivalence interval
- data: A description of the data
 - type: The type of data ('raw' when arguments x and y are used or 'summary' when arguments n_x, n_y, mean_x, mean_y, sd_x, and sd_y (or ci_margin and ci_level instead of sd_x and sd_y) are used)
 - ...: values for the arguments used, depending on 'raw' or 'summary'
- prior_scale: The width of the Cauchy prior distribution
- bf: The resulting Bayes factor

A summary of the model is shown by printing the object.

References

- Gronau, Q. F., Ly, A., & Wagenmakers, E.-J. (2020). Informed Bayesian t-tests. *The American Statistician*, 74, 137-143.
- Liang, F., Paulo, R., Molina, G., Clyde, M. A., & Berger, J. O. (2008). Mixtures of g priors for Bayesian variable selection. *Journal of the American Statistical Association*, 103, 410-423.
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16, 225-237.
- van Ravenzwaaij, D., Monden, R., Tendeiro, J. N., & Ioannidis, J. P. A. (2019). Bayes factors for superiority, non-inferiority, and equivalence designs. *BMC Medical Research Methodology*, 19, 71.

Examples

```
## equiv_bf using raw data:

# Assign model to variable.
equiv_raw <- equiv_bf(x = rnorm(100, 10, 15),
                     y = rnorm(130, 13, 10))
```

```
# Extract Bayes factor from variable.
get_bf(equiv_raw)

# -----
# -----

## equiv_bf using summary statistics with data from Steiner et al. (2015).
## With a point null hypothesis:

# Assign model to variable.
equiv_sum_point <- equiv_bf(n_x = 560,
                           n_y = 538,
                           mean_x = 8.683,
                           mean_y = 8.516,
                           sd_x = 3.6,
                           sd_y = 3.6)

# Extract Bayes factor from model.
get_bf(equiv_sum_point)

# -----
# -----

## equiv_bf using summary statistics with data from Steiner et al. (2015).
## With an interval null hypothesis:

# Assign model to variable.
equiv_sum_interval <- equiv_bf(n_x = 560,
                              n_y = 538,
                              mean_x = 8.683,
                              mean_y = 8.516,
                              sd_x = 3.6,
                              sd_y = 3.6,
                              interval = 0.05)

# Extract Bayes factor from model.
get_bf(equiv_sum_interval)
```

get_bf

Extract Bayes factor from S4 object

Description

`get_bf` extracts the Bayes factor from an S4 object (i.e., `baymedrSuperiority`, `baymedrEquivalence`, `baymedrNonInferiority`), `baymedrCoxProportionalHazards`, and `baymedrCoxProportionalHazards-Multi`.

Usage

```
get_bf(object)
```

Arguments

object An S4 object of class `baymedrSuperiority`, `baymedrEquivalence`, `baymedrNonInferiority`, `baymedrCoxProportionalHazards`, or `baymedrCoxProportionalHazardsMulti`.

Value

A numeric vector, providing the Bayes factor(s) from an S4 object.

Examples

```
# Extract Bayes factor from a baymedrSuperiority object using raw data:
mod_super <- super_bf(x = rnorm(100, 10, 15),
                     y = rnorm(130, 13, 10))

get_bf(object = mod_super)

# Extract Bayes factor from a baymedrEquivalence object using raw data:
mod_equiv <- equiv_bf(x = rnorm(100, 10, 15),
                     y = rnorm(130, 13, 10))

get_bf(object = mod_equiv)

# Extract Bayes factor from a baymedrNonInferiority object using raw data:
mod_infer <- infer_bf(x = rnorm(100, 10, 15),
                     y = rnorm(130, 13, 10),
                     ni_margin = 1)

get_bf(object = mod_infer)

# Extract Bayes factor from a baymedrCoxProportionalHazards object:
data <- survival::aml
names(data) <- c("time", "event", "group")
data$group <- ifelse(test = data$group == "Maintained",
                    yes = 0,
                    no = 1)

mod_coxph <- coxph_bf(data = data)

get_bf(object = mod_coxph)
```

infer_bf

Bayes factor for non-inferiority designs

Description

`infer_bf` computes a Bayes factor for non-inferiority designs with a continuous dependent variable.

Usage

```
infer_bf(
  x = NULL,
  y = NULL,
  n_x = NULL,
  n_y = NULL,
  mean_x = NULL,
  mean_y = NULL,
  sd_x = NULL,
  sd_y = NULL,
  ci_margin = NULL,
  ci_level = NULL,
  ni_margin = NULL,
  ni_margin_std = TRUE,
  prior_scale = 1/sqrt(2),
  direction = "high"
)
```

Arguments

x	A numeric vector of observations for the control group.
y	A numeric vector of observations for the experimental group.
n_x	A numeric vector of length one, specifying the sample size of the control group.
n_y	A numeric vector of length one, specifying the sample size of the experimental group.
mean_x	A numeric vector of length one, specifying the mean of the dependent variable in the control group.
mean_y	A numeric vector of length one, specifying the mean of the dependent variable in the experimental group.
sd_x	A numeric vector of length one, specifying the standard deviation of the dependent variable in the control group. Only sd_x and sd_y OR ci_margin and ci_level should be defined (see Details).
sd_y	A numeric vector of length one, specifying the standard deviation of the dependent variable in the experimental group. Only sd_x and sd_y OR ci_margin and ci_level should be defined (see Details).
ci_margin	A numeric vector of length one, specifying the margin of the confidence interval (i.e., the width of the confidence interval divided by 2) of the mean difference on the dependent variable between the experimental and control groups. The value should be a positive number Only sd_x and sd_y OR ci_margin and ci_level should be defined (see Details).
ci_level	A numeric vector of length one, specifying the confidence level of ci_margin. The value must be between 0 and 1 (e.g., 0.95 for a 95% confidence interval). Only sd_x and sd_y OR ci_margin and ci_level should be defined (see Details).
ni_margin	A numeric vector of length one, specifying the non-inferiority margin. The value should be a positive number.

<code>ni_margin_std</code>	A logical vector of length one, specifying whether the non-inferiority margin (i.e., <code>ni_margin</code>) is given in standardized (TRUE; the default) or unstandardized (FALSE) units.
<code>prior_scale</code>	A numeric vector of length one, specifying the scale of the Cauchy prior distribution for the effect size under the alternative hypothesis (see Details). The default value is $r = 1 / \sqrt{2}$.
<code>direction</code>	A character vector of length one, specifying the direction of non-inferior scores. 'low' indicates that low scores on the measure of interest correspond to a non-inferior outcome and 'high' (the default) indicates that high scores on the measure of interest correspond to a non-inferior outcome (see Details).

Details

The formulation of the null and alternative hypotheses for the non-inferiority design differs depending on whether high or low scores on the dependent variable represent non-inferiority. In the case where high scores correspond to non-inferiority, the hypotheses are as follows: The null hypothesis states that the population mean of the experimental group (e.g., a new medication) is lower than the population mean of the control group (e.g., a placebo or an already existing medication) minus the non-inferiority margin. The alternative hypothesis states that the population mean of the experimental group is higher than the population mean of the control group minus the non-inferiority margin. Thus, the null hypothesis goes in the negative direction (i.e., H_-) and the alternative hypothesis in the positive direction (i.e., H_+). In turn, in the case where low scores correspond to non-inferiority, the hypotheses are as follows: The null hypothesis states that the population mean of the experimental group is higher than the population mean of the control group plus the non-inferiority margin. The alternative hypothesis states that the population mean of the experimental group is lower than the population mean of the control group plus the non-inferiority margin. Thus, the null hypothesis goes in the positive direction (i.e., H_+) and the alternative hypothesis in the negative direction (i.e., H_-). The dependent variable must be continuous.

Since the main goal of `infer_bf` is to establish non-inferiority, the resulting Bayes factor quantifies evidence in favor of the alternative hypothesis. In the case where high values represent non-inferiority we have BF_{+-} and in the case where low values represent non-inferiority we have BF_{-+} . Evidence for the null hypothesis can easily be calculated by taking the reciprocal of the original Bayes factor (i.e., $BF_{+-} = 1 / BF_{-+}$ and vice versa). Quantification of evidence in favor of the null hypothesis is logically sound and legitimate within the Bayesian framework (see e.g., van Ravenzwaaij et al., 2019).

`infer_bf` can be utilized to calculate a Bayes factor based on raw data (i.e., if arguments `x` and `y` are defined) or summary statistics (i.e., if arguments `n_x`, `n_y`, `mean_x`, and `mean_y` (or `ci_margin` and `ci_level` instead of `sd_x` and `sd_y`) are defined). Arguments with 'x' as a name or suffix correspond to the control group, whereas arguments with 'y' as a name or suffix correspond to the experimental group.

Since sometimes high scores on the dependent variable are considered non-inferior (e.g., amount of social interactions) and sometimes rather the low scores (e.g., severity of symptoms), the direction of non-inferiority can be specified with the argument `direction`. For the case where high values on the dependent variable indicate non-inferiority, 'high' (the default) should be specified for the argument `direction`; if low values on the dependent variable indicate non-inferiority, 'low' should be specified for the argument `direction`.

With the argument `ni_margin`, the non-inferiority margin can be specified. `ni_margin` should be a positive number.' It can be declared whether the non-inferiority margin is specified in standardized or unstandardized units with the `ni_margin_std` argument, where `TRUE`, corresponding to standardized units, is the default.

For the calculation of the Bayes factor, a Cauchy prior density centered on 0 is chosen for the effect size under the alternative hypothesis. The standard Cauchy distribution, with a location parameter of 0 and a scale parameter of 1, resembles a standard Normal distribution, except that the Cauchy distribution has less mass at the center but heavier tails (Liang et al., 2008; Rouder et al., 2009). The argument `prior_scale` specifies the width of the Cauchy prior, which corresponds to half of the interquartile range. Thus, by adjusting the Cauchy prior scale with `prior_scale`, different ranges of expected effect sizes can be emphasized. The default prior scale is set to $r = 1 / \sqrt{2}$.

`infer_bf` creates an S4 object of class `baymedrNonInferiority`, which has multiple slots/entries (e.g., type of data, prior scale, Bayes factor, etc.; see Value). If it is desired to store or extract solely the Bayes factor, the user can do this with `get_bf`, by setting the S4 object as an argument (see Examples).

Value

An S4 object of class `baymedrNonInferiority` is returned. Contained are a description of the model and the resulting Bayes factor:

- test: The type of analysis
- hypotheses: A statement of the hypotheses
 - h0: The null hypothesis
 - h1: The alternative hypothesis
- ni_margin: The value for `ni_margin` in standardized and unstandardized units
 - `ni_mar_std`: The standardized non-inferiority margin
 - `ni_mar_unstd`: The unstandardized non-inferiority margin
- data: A description of the data
 - type: The type of data ('raw' when arguments `x` and `y` are used or 'summary' when arguments `n_x`, `n_y`, `mean_x`, `mean_y`, `sd_x`, and `sd_y` (or `ci_margin` and `ci_level` instead of `sd_x` and `sd_y`) are used)
 - ...: values for the arguments used, depending on 'raw' or 'summary'
- prior_scale: The width of the Cauchy prior distribution
- bf: The resulting Bayes factor

A summary of the model is shown by printing the object.

References

- Gronau, Q. F., Ly, A., & Wagenmakers, E.-J. (2020). Informed Bayesian t-tests. *The American Statistician*, 74, 137-143.
- Liang, F., Paulo, R., Molina, G., Clyde, M. A., & Berger, J. O. (2008). Mixtures of g priors for Bayesian variable selection. *Journal of the American Statistical Association*, 103, 410-423.
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16, 225-237.

van Ravenzwaaij, D., Monden, R., Tendeiro, J. N., & Ioannidis, J. P. A. (2019). Bayes factors for superiority, non-inferiority, and equivalence designs. *BMC Medical Research Methodology*, 19, 71.

Examples

```
## infer_bf using raw data:

# Assign model to variable.
infer_raw <- infer_bf(x = rnorm(100, 10, 15),
                     y = rnorm(130, 13, 10),
                     ni_margin = 1.5,
                     ni_margin_std = FALSE)

# Extract Bayes factor from model.
get_bf(infer_raw)

# -----
# -----

## infer_bf using summary statistics with data from Andersson et al. (2013).
## Test at timepoint 1:

# Assign model to variable.
infer_sum_t1 <- infer_bf(n_x = 33,
                        n_y = 32,
                        mean_x = 17.1,
                        mean_y = 13.6,
                        sd_x = 8,
                        sd_y = 9.8,
                        ni_margin = 2,
                        ni_margin_std = FALSE,
                        direction = "low")

# Extract Bayes factor from model
get_bf(infer_sum_t1)

# -----
# -----

## infer_bf using summary statistics with data from Andersson et al. (2013).
## Test at timepoint 2:

# Assign model to variable.
infer_sum_t2 <- infer_bf(n_x = 30,
                        n_y = 32,
                        mean_x = 13.5,
                        mean_y = 9.2,
                        sd_x = 8.7,
                        sd_y = 7.6,
                        ni_margin = 2,
                        ni_margin_std = FALSE,
                        direction = "low")
```

```
# Extract Bayes factor from model
get_bf(infer_sum_t2)
```

model-classes	<i>S4 classes to represent different models</i>
---------------	---

Description

The S4 classes `baymedrSuperiority`, `baymedrEquivalence`, `baymedrNonInferiority`, `baymedrCoxProportionalHazards`, and `baymedrCoxProportionalHazardsMulti` represent models for the superiority (`super_bf`), equivalence (`equiv_bf`), non-inferiority (`infer_bf`), and Cox proportional hazards (`coxph_bf`) models, respectively.

Slots

`test` Type of test that was conducted.
`hypotheses` The hypotheses that are tested.
`data` The type of data that was used.
`prior_scale` The Cauchy prior scale that was used.
`bf` The resulting Bayes factor.
`interval` The equivalence interval in case of `equiv_bf`.
`ni_margin` The non-inferiority margin in case of `infer_bf`.
`prior` The mean and standard deviation of the Normal prior in case of `coxph_bf`.

<code>super_bf</code>	<i>Bayes factor for superiority designs</i>
-----------------------	---

Description

`super_bf` computes a Bayes factor for superiority designs with a continuous dependent variable.

Usage

```
super_bf(
  x = NULL,
  y = NULL,
  n_x = NULL,
  n_y = NULL,
  mean_x = NULL,
  mean_y = NULL,
  sd_x = NULL,
  sd_y = NULL,
  ci_margin = NULL,
  ci_level = NULL,
  prior_scale = 1/sqrt(2),
  direction = "high"
)
```

Arguments

x	A numeric vector of observations for the control group.
y	A numeric vector of observations for the experimental group.
n_x	A numeric vector of length one, specifying the sample size of the control group.
n_y	A numeric vector of length one, specifying the sample size of the experimental group.
mean_x	A numeric vector of length one, specifying the mean of the dependent variable in the control group.
mean_y	A numeric vector of length one, specifying the mean of the dependent variable in the experimental group.
sd_x	A numeric vector of length one, specifying the standard deviation of the dependent variable in the control group. Only sd_x and sd_y OR ci_margin and ci_level should be defined (see Details).
sd_y	A numeric vector of length one, specifying the standard deviation of the dependent variable in the experimental group. Only sd_x and sd_y OR ci_margin and ci_level should be defined (see Details).
ci_margin	A numeric vector of length one, specifying the margin of the confidence interval (i.e., the width of the confidence interval divided by 2) of the mean difference on the dependent variable between the experimental and control groups. The value should be a positive number Only sd_x and sd_y OR ci_margin and ci_level should be defined (see Details).
ci_level	A numeric vector of length one, specifying the confidence level of ci_margin. The value must be between 0 and 1 (e.g., 0.95 for a 95% confidence interval). Only sd_x and sd_y OR ci_margin and ci_level should be defined (see Details).
prior_scale	A numeric vector of length one, specifying the scale of the Cauchy prior distribution for the effect size under the alternative hypothesis (see Details). The default value is $r = 1 / \sqrt{2}$.
direction	A character vector of length one, specifying the direction of superior scores. 'low' indicates that low scores on the measure of interest correspond to a superior outcome and 'high' (the default) indicates that high scores on the measure of interest correspond to a superior outcome (see Details).

Details

The formulation of the null and alternative hypotheses for the superiority design differs depending on whether high or low scores on the dependent variable represent superiority. In both cases, the null hypothesis (i.e., H_0) states that the population means of the experimental group and the control group are equivalent. In the case where high scores correspond to superiority, the alternative hypothesis states that the population mean of the experimental group is higher than the population mean of the control group. Thus, the alternative hypothesis goes in the positive direction (i.e., H_+). In turn, in the case where low scores correspond to superiority, the alternative hypothesis states that the population mean of the experimental group is lower than the population mean of the control group. Thus, the alternative hypothesis goes in the negative direction (i.e., H_-). The dependent variable must be continuous.

Since the main goal of `super_bf` is to establish superiority, the resulting Bayes factor quantifies evidence in favor of the alternative hypothesis. In the case where low values represent superiority we have BF-0, whereas in the case where high values represent superiority we have BF+0. Evidence for the null hypothesis can easily be calculated by taking the reciprocal of the original Bayes factor (i.e., $BF_{0-} = 1 / BF_{-0}$ and $BF_{0+} = 1 / BF_{+0}$). Quantification of evidence in favor of the null hypothesis is logically sound and legitimate within the Bayesian framework (see e.g., van Ravenzwaaij et al., 2019).

`super_bf` can be utilized to calculate a Bayes factor based on raw data (i.e., if arguments `x` and `y` are defined) or summary statistics (i.e., if arguments `n_x`, `n_y`, `mean_x`, and `mean_y` are defined). In the latter case, the user has the freedom to supply values either for the arguments `sd_x` and `sd_y` **OR** `ci_margin` and `ci_level`. Arguments with 'x' as a name or suffix correspond to the control group, whereas arguments with 'y' as a name or suffix correspond to the experimental group (i.e., the group for which we seek to establish superiority).

For the calculation of the Bayes factor, a Cauchy prior density centered on 0 is chosen for the effect size under the alternative hypothesis. The standard Cauchy distribution, with a location parameter of 0 and a scale parameter of 1, resembles a standard Normal distribution, except that the Cauchy distribution has less mass at the center but heavier tails (Liang et al., 2008; Rouder et al., 2009). The argument `prior_scale` specifies the width of the Cauchy prior, which corresponds to half of the interquartile range. Thus, by adjusting the Cauchy prior scale with `prior_scale`, different ranges of expected effect sizes can be emphasized. The default prior scale is set to $r = 1 / \sqrt{2}$.

`super_bf` creates an S4 object of class `baymedrSuperiority`, which has multiple slots/entries (e.g., type of data, prior scale, Bayes factor, etc.; see Value). If it is desired to store or extract solely the Bayes factor, the user can do this with `get_bf`, by setting the S4 object as an argument (see Examples).

Value

An S4 object of class `baymedrSuperiority` is returned. Contained are a description of the model and the resulting Bayes factor:

- test: The type of analysis
- hypotheses: A statement of the hypotheses
 - h0: The null hypothesis
 - h1: The alternative hypothesis
- data: A description of the data
 - type: The type of data ('raw' when arguments `x` and `y` are used or 'summary' when arguments `n_x`, `n_y`, `mean_x`, `mean_y`, `sd_x`, and `sd_y` (or `ci_margin` and `ci_level` instead of `sd_x` and `sd_y`) are used)
 - ...: values for the arguments used, depending on 'raw' or 'summary'
- prior_scale: The scale of the Cauchy prior distribution
- bf: The resulting Bayes factor

A summary of the model is shown by printing the object.


```
ci_level = 0.95,  
direction = "low")  
  
# Extract Bayes factor from model.  
get_bf(super_sum_ex2)
```

Index

baymedr (baymedr-package), [2](#)
baymedr-package, [2](#)
baymedrCoxProportionalHazards, [4](#), [11](#), [12](#),
[17](#)
baymedrCoxProportionalHazards-class
 (model-classes), [17](#)
baymedrCoxProportionalHazardsMulti, [11](#),
[12](#), [17](#)
baymedrCoxProportionalHazardsMulti-class
 (model-classes), [17](#)
baymedrEquivalence, [9–12](#), [17](#)
baymedrEquivalence-class
 (model-classes), [17](#)
baymedrNonInferiority, [11](#), [12](#), [15](#), [17](#)
baymedrNonInferiority-class
 (model-classes), [17](#)
baymedrSuperiority, [11](#), [12](#), [17](#), [19](#)
baymedrSuperiority-class
 (model-classes), [17](#)

coxph_bf, [3](#), [3](#), [4](#), [7](#), [17](#)
coxph_data_sim, [4](#), [5](#), [5](#), [6](#)

equiv_bf, [7](#), [7](#), [9](#), [17](#)

get_bf, [4](#), [10](#), [11](#), [11](#), [15](#), [19](#)

infer_bf, [12](#), [12](#), [14](#), [15](#), [17](#)

model-classes, [17](#)

psoptim, [6](#), [7](#)

super_bf, [17](#), [17](#), [19](#)