Package: simFastBOIN (via r-universe)

October 16, 2026

Type Package

Title Fast Simulation of Bayesian Optimal Interval Designs for Phase I Trials

Version 2.1.0

Date 2026-10-13

Description Design and evaluate phase I dose-finding trials that use the Bayesian optimal interval (BOIN) design of Liu and Yuan (2015) <doi:10.1111/rssc.12089>. Functions are provided to tabulate the decision boundaries, to simulate trials, to estimate the dose-toxicity curve under a monotonicity constraint and to select the maximum tolerated dose. The simulation engine is written in C++ and draws one random variate per patient in enrollment order, which reproduces the reference implementation in the 'BOIN' package trial by trial for a given seed. The traditional 3+3 design is provided as a comparator, with operating characteristics obtained in closed form rather than by simulation. For toxicity that is observed late, the time-to-event BOIN design is provided with the two estimators of Yuan and colleagues (2018) <doi:10.1158/1078-0432.CCR-18-0246> and Lin and Yuan (2020) <doi:10.1093/biostatistics/kxz007>. Patients who arrive while the dose escalation waits can be backfilled to lower doses with the designs of Zhao and colleagues (2024) <doi:10.1158/1078-0432.CCR-23-2585> and Chen and colleagues (2026) <doi:10.1007/s43441-026-00994-0>, the latter with the suspension rules of Chen and colleagues (2025) <doi:10.1158/1078-0432.CCR-24-3669>. Trials of two drugs given together are handled by the drug-combination design of Lin and Yin (2017) <doi:10.1177/0962280215594494>.

License MIT + file LICENSE

Encoding UTF-8

Depends R (>= 3.5.0)

Imports Iso, Rcpp (>= 1.0.0), stats, utils

LinkingTo Rcpp

Suggests BOIN, ggplot2 (>= 3.4.0), knitr, rmarkdown, spelling, testthat (>= 3.0.0)

Language en-US

URL https://gosukehommaex.github.io/simFastBOIN/, https://github.com/gosukehommaEX/simFastBOIN

BugReports https://github.com/gosukehommaEX/simFastBOIN/issues

Config/testthat/edition 3

VignetteBuilder knitr

Config/roxygen2/version 8.0.0

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

Date/Publication 2026-10-16 09:43:08 UTC

RemoteUrl https://github.com/gosukehommaex/simfastboin

RemoteRef HEAD

RemoteSha 7921ffaba7b50b2e0ec919f9831f7d46d52a88b6

Contents

as.data.frame.boin_boundary . . . 3
be_boin_simulate . . . 4
bf_boin_simulate . . . 8
boin_boundary . . . 13
boin_decision_table . . . 15
boin_isotonic . . . 16
boin_lambda . . . 18
boin_p_tox . . . 19
boin_select_mtd . . . 20
boin_simulate . . . 22
boin_stopping_table . . . 24
comb_boin_next . . . 25
comb_boin_select_mtd . . . 28
comb_boin_simulate . . . 30
decision_table_3p3 . . . 33
oc_3p3 . . . 34
plot.boin_decision_table . . . 36
plot.comb_boin_next . . . 37
plot.decision_table_3p3 . . . 38
plot.tite_boin_decision_table . . . 39
print.backfill_oc . . . 41
print.backfill_oc_multi . . . 42
print.backfill_trials . . . 43
print.boin_boundary . . . 44
print.boin_decision_table . . . 45
print.boin_oc . . . 46

print.boin_oc_multi 47
print.boin_trials 48
print.comb_boin_next 49
print.comb_boin_oc 49
print.comb_boin_oc_multi 50
print.comb_boin_trials 51
print.decision_table_3p3 52
print.oc_3p3 53
print.tite_boin_decision_table 54
print.tite_boin_oc 55
print.tite_boin_oc_multi 56
print.tite_boin_trials 58
sim_3p3 58
sim_be_boin 60
sim_be_boin_multi 64
sim_bf_boin 67
sim_bf_boin_multi 72
sim_boin 76
sim_boin_multi 79
sim_comb_boin 82
sim_comb_boin_multi 85
sim_tite_boin 87
sim_tite_boin_multi 91
tite_boin_decision_table 95
tite_boin_simulate 98

Index **103**

`as.data.frame.boin_boundary` *Coerce BOIN Decision Boundaries to a Data Frame*

Description

Return the integer decision boundaries produced by `boin_boundary` as one row per sample size.

Usage

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

Arguments

`x`	An object of class `boin_boundary`.
`row.names`	Passed on for compatibility with the generic, normally unused.
`optional`	Passed on for compatibility with the generic, normally unused.
`...`	Further arguments, currently ignored.

Value

A data frame with columns `n_pts`, `escalate_if_dlt_leq`, `deescalate_if_dlt_geq` and `eliminate_if_dlt_geq`, plus `stop_if_dlt_geq` when the boundaries were built with `extrasafe = TRUE`.

Examples

```
bd <- boin_boundary(target = 0.30, max_n = 18)
as.data.frame(bd)[seq(3, 18, by = 3), ]
```

`be_boin_simulate` *Simulate BE-BOIN Trials*

Description

Run the BOIN design with backfilling and late-onset toxicity (BE-BOIN) of Chen et al. (2026) over many simulated trials. Dose escalation follows TITE-BOIN, which decides while some patients are still pending; patients who arrive while the escalation is suspended are backfilled to a lower dose that is safe and has shown a response. Every estimate imputes the pending patients, backfilled ones included. Return the patient, DLT, backfill and response counts at every dose together with the duration of each trial.

Usage

```
be_boin_simulate(
  target,
  p_true,
  p_resp,
  n_cohort,
  cohort_size,
  window,
  accrual_rate,
  n_cap = 12,
  backfill_dose = c("highest", "lowest"),
  conflict_dose = c("lowest", "highest"),
  no_slot = c("wait", "leave"),
  accrual = c("exponential", "uniform", "fixed"),
  dlt_time = c("weibull", "uniform"),
  late_fraction = 0.5,
  resp_window = window,
  resp_late_fraction = 0.5,
  resp_cor = 0,
  max_pending_ratio = 0.49,
  min_completed = 0,
  min_follow_up = 0.25,
  n_trials = 10000,
  start_dose = 1,
```

```
  n_earlystop = 18,
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  n_earlystop_rule = c("with_stay", "simple"),
  seed = 123
)
```

Arguments

`target`	Numeric scalar. Target DLT probability, for example 0.25.
`p_true`	Numeric vector. True DLT probability at each dose level, in increasing dose order. The probability of a DLT within the assessment window.
`p_resp`	Numeric vector of the same length as `p_true`. True probability of a response within the response window at each dose. A response at or below a dose is what makes the dose eligible for backfilling.
`n_cohort`	Integer scalar. Number of cohorts in the dose escalation.
`cohort_size`	Integer scalar or vector. Number of patients per cohort of the dose escalation. A scalar is used for every cohort. A vector shorter than `n_cohort` is padded with its last element and a longer one is truncated.
`window`	Numeric scalar. Length of the DLT assessment window.
`accrual_rate`	Numeric scalar. Average number of patients arriving per unit of time, in the same unit as `window`.
`n_cap`	Integer scalar. A dose is no longer backfilled once this many patients, from the dose escalation and from backfilling together, have been treated at it. The dose escalation can still treat patients at the dose. Defaults to 12.
`backfill_dose`	Character scalar. Which dose receives a backfill patient when several are open: `"highest"` (the default, as in both articles) or `"lowest"`.
`conflict_dose`	Character scalar. When the data of several backfilled doses conflict with those of the current dose, which of them anchors the pooled estimate: `"lowest"` (the default here, following Chen et al., 2026) or `"highest"` (the default of `bf_boin_simulate`, following Zhao et al., 2024).
`no_slot`	Character scalar. What happens to a patient who arrives when neither the dose escalation nor any backfill dose can take a patient: `"wait"` (the default), the patient waits and is treated as soon as one can, as in `tite_boin_simulate`; or `"leave"`, the patient is turned away and the trial waits for the next arrival.
`accrual`	Character scalar. Distribution of the times between arrivals: `"exponential"` (the default) for a Poisson process, `"uniform"` for a uniform distribution on `(0, 2 / accrual_rate)`, or `"fixed"` for arrivals exactly `1 / accrual_rate` apart.
`dlt_time`	Character scalar. Distribution of the time to DLT within the window, and of the time to response within the response window: `"weibull"` (the default) or `"uniform"`.

late_fraction
: Numeric scalar strictly between 0 and 1. Share of the DLTs that occur in the second half of the assessment window under the Weibull distribution. Defaults to 0.5.

resp_window
: Numeric scalar. Length of the window within which a response is observed. Defaults to `window`.

resp_late_fraction
: Numeric scalar strictly between 0 and 1. Share of the responses that occur in the second half of the response window under the Weibull distribution. Defaults to 0.5.

resp_cor
: Numeric scalar strictly between -1 and 1. Correlation of the Gaussian copula that links the DLT and the response of a patient; a positive value makes a response more likely in a patient with a DLT. Defaults to 0, independence.

max_pending_ratio
: Numeric scalar greater than 0 and at most 1. Rule 1 of the design: the escalation is suspended when the proportion of pending patients at the current dose exceeds this value, unless the dose is de-escalated whatever their outcomes. Defaults to 0.49, which suspends when fewer than 51 percent have completed the assessment, as in Chen et al. (2026). See `tite_boin_decision_table`.

min_completed
: Integer scalar. While some patients are pending, escalation requires at least this many patients with a completed assessment. Defaults to 0, which the design does not use. See `tite_boin_decision_table`.

min_follow_up
: Numeric scalar between 0 and 1. Rule 2 of the design: escalation requires every pending patient at the current dose to have been followed for at least this fraction of the window. Defaults to 0.25, as in Chen et al. (2026). See `tite_boin_decision_table`.

n_trials
: Integer scalar. Number of trials to simulate. Defaults to 10000.

start_dose
: Integer scalar. Dose level for the first cohort. Defaults to 1.

n_earlystop
: Integer scalar. The trial stops once this many patients, backfilled ones included, have been treated at the current dose and the design would stay there. Defaults to 18.

p_saf
: Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to `0.6 * target`.

p_tox
: Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to `1.4 * target`.

cutoff_eli
: Numeric scalar. Posterior probability cutoff for dose elimination. Defaults to 0.95.

extrasafe
: Logical scalar. Apply the stricter safety stopping rule at the lowest dose. Defaults to `FALSE`.

offset
: Numeric scalar between 0 and 0.5. Amount by which `cutoff_eli` is relaxed for the safety stopping rule. Defaults to 0.05.

n_earlystop_rule
: Character scalar, either `"with_stay"` or `"simple"`. See `boin_simulate`.

seed
: Integer scalar or `NULL`. Random seed. When `NULL` the current state of the random number generator is used. The state of the calling session is restored on exit. Defaults to 123.

Details

Dose escalation. The dose for a cohort is decided when its first patient arrives, with the rules of `tite_boin_decision_table` for `method = "imputation"` and the suspension rules `max_pending_ratio` and `min_follow_up`; every patient at the current dose counts, backfilled ones included. A suspension lasts until a decision can be taken again, as in `tite_boin_simulate`. Elimination counts every treated patient, the pending ones as without DLT, and is checked at every dose.

Estimates. The DLT rate at a dose is the single mean imputation estimate of TITE-BOIN, equation (1) of Chen et al. (2026), and the rate pooled over several doses adds up the observed and imputed DLTs of each dose, each imputed with its own posterior mean, and divides by the number of treated patients, equation (2). Both enter the opening and closing of backfill doses and the resolution of conflicts, which otherwise follow `bf_boin_simulate`. The categories of a backfilled dose for the conflicts compare its estimate with the escalation and de-escalation boundaries directly.

End of the trial, arrivals and random numbers are as in `bf_boin_simulate`. When no dose is ever opened for backfilling, for example with every `p_resp` equal to 0, and with `no_slot = "wait"`, the trials are identical to those of `tite_boin_simulate` with `method = "imputation"`, the same suspension rules and the same seed.

Chen et al. (2026) report their operating characteristics for scenarios that are given only as figures, so they cannot be reproduced. The decision table of their supplementary materials is reproduced by `tite_boin_decision_table`, and their trial example is replayed patient by patient in the tests of the package.

Value

An object of class `c("be_boin_trials", "backfill_trials")`, with the components described in `bf_boin_simulate`.

References

Chen, K., Zhao, Y., Takeda, K. and Yuan, Y. (2026). BE-BOIN: A Dose Optimization Design Accommodating Backfill and Late-Onset Toxicity. Therapeutic Innovation and Regulatory Science. doi:10.1007/s43441026009940

Chen, K., Chen, T.-Y., Zhang, Y., Lin, R. and Yuan, Y. (2025). Practical Considerations for Using the TITE-BOIN Design to Handle Late-Onset Toxicity or Fast Accrual in Phase I Trials. Clinical Cancer Research, 31(13), 2573-2580.

Zhao, Y., Yuan, Y., Korn, E. L. and Freidlin, B. (2024). Backfilling Patients in Phase I Dose-Escalation Trials Using Bayesian Optimal Interval Design (BOIN). Clinical Cancer Research, 30(4), 673-679.

See Also

`bf_boin_simulate`, `tite_boin_simulate`

Examples

```
# A three month window with two patients a month
trials <- be_boin_simulate(
  target = 0.25,
  p_true = c(0.04, 0.12, 0.25, 0.43, 0.63),
```

```
  p_resp = c(0.10, 0.20, 0.30, 0.45, 0.58),
  n_cohort = 10,
  cohort_size = 3,
  window = 3,
  accrual_rate = 2,
  n_earlystop = 9,
  n_trials = 200,
  seed = 123
)
trials

summary(trials$duration)
```

`bf_boin_simulate` *Simulate BF-BOIN Trials*

Description

Run the BOIN design with backfilling (BF-BOIN) of Zhao et al. (2024) over many simulated trials. Dose escalation proceeds in fully staggered cohorts with the BOIN decision table; patients who arrive while the escalation waits for the DLT assessment of its current cohort are backfilled to a lower dose that is safe and has shown a response. Return the patient, DLT, backfill and response counts at every dose together with the duration of each trial.

Usage

```
bf_boin_simulate(
  target,
  p_true,
  p_resp,
  n_cohort,
  cohort_size,
  window,
  accrual_rate,
  n_cap = 12,
  backfill_dose = c("highest", "lowest"),
  conflict_dose = c("highest", "lowest"),
  no_slot = c("wait", "leave"),
  accrual = c("exponential", "uniform", "fixed"),
  dlt_time = c("weibull", "uniform"),
  late_fraction = 0.5,
  resp_window = window,
  resp_late_fraction = 0.5,
  resp_cor = 0,
  n_trials = 10000,
  start_dose = 1,
  n_earlystop = 18,
```

```
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  stay_on_1_of_3 = FALSE,
  n_earlystop_rule = c("with_stay", "simple"),
  seed = 123
)
```

Arguments

`target`	Numeric scalar. Target DLT probability, for example 0.25.
`p_true`	Numeric vector. True DLT probability at each dose level, in increasing dose order. The probability of a DLT within the assessment window.
`p_resp`	Numeric vector of the same length as `p_true`. True probability of a response within the response window at each dose. A response at or below a dose is what makes the dose eligible for backfilling.
`n_cohort`	Integer scalar. Number of cohorts in the dose escalation.
`cohort_size`	Integer scalar or vector. Number of patients per cohort of the dose escalation. A scalar is used for every cohort. A vector shorter than `n_cohort` is padded with its last element and a longer one is truncated.
`window`	Numeric scalar. Length of the DLT assessment window.
`accrual_rate`	Numeric scalar. Average number of patients arriving per unit of time, in the same unit as `window`.
`n_cap`	Integer scalar. A dose is no longer backfilled once this many patients, from the dose escalation and from backfilling together, have been treated at it. The dose escalation can still treat patients at the dose. Defaults to 12.
`backfill_dose`	Character scalar. Which dose receives a backfill patient when several are open: `"highest"` (the default, as in both articles) or `"lowest"`.
`conflict_dose`	Character scalar. When the data of several backfilled doses conflict with those of the current dose, which of them anchors the pooled estimate: `"highest"` (the default here, following Zhao et al., 2024) or `"lowest"` (the default of `be_boin_simulate`, following Chen et al., 2026). See the details.
`no_slot`	Character scalar. What happens to a patient who arrives when neither the dose escalation nor any backfill dose can take a patient: `"wait"` (the default), the patient waits and is treated as soon as one can, as in `tite_boin_simulate`; or `"leave"`, the patient is turned away and the trial waits for the next arrival.
`accrual`	Character scalar. Distribution of the times between arrivals: `"exponential"` (the default) for a Poisson process, `"uniform"` for a uniform distribution on `(0, 2 / accrual_rate)`, or `"fixed"` for arrivals exactly `1 / accrual_rate` apart.
`dlt_time`	Character scalar. Distribution of the time to DLT within the window, and of the time to response within the response window: `"weibull"` (the default) or `"uniform"`.

late_fraction Numeric scalar strictly between 0 and 1. Share of the DLTs that occur in the second half of the assessment window under the Weibull distribution. Defaults to 0.5.

resp_window Numeric scalar. Length of the window within which a response is observed. Defaults to window.

resp_late_fraction
Numeric scalar strictly between 0 and 1. Share of the responses that occur in the second half of the response window under the Weibull distribution. Defaults to 0.5.

resp_cor Numeric scalar strictly between -1 and 1. Correlation of the Gaussian copula that links the DLT and the response of a patient; a positive value makes a response more likely in a patient with a DLT. Defaults to 0, independence.

n_trials Integer scalar. Number of trials to simulate. Defaults to 10000.

start_dose Integer scalar. Dose level for the first cohort. Defaults to 1.

n_earlystop Integer scalar. The trial stops once this many patients, backfilled ones included, have been treated at the current dose and the design would stay there. Defaults to 18.

p_saf Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to 0.6 * target.

p_tox Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to 1.4 * target.

cutoff_eli Numeric scalar. Posterior probability cutoff for dose elimination. Defaults to 0.95.

extrasafe Logical scalar. Apply the stricter safety stopping rule at the lowest dose. Defaults to FALSE.

offset Numeric scalar between 0 and 0.5. Amount by which cutoff_eli is relaxed for the safety stopping rule. Defaults to 0.05.

stay_on_1_of_3 Logical scalar. When TRUE, one DLT out of three patients leads to staying rather than de-escalating, in every use of the decision table. Zhao et al. (2024) use this modification in their main simulation. Defaults to FALSE. See boin_boundary.

n_earlystop_rule
Character scalar, either "with_stay" or "simple". See boin_simulate.

seed Integer scalar or NULL. Random seed. When NULL the current state of the random number generator is used. The state of the calling session is restored on exit. Defaults to 123.

Details

Dose escalation. Cohorts are fully staggered: the next cohort starts only when every patient of the current one has completed the DLT assessment. The decision then follows the BOIN decision table applied to the patients at the current dose who have completed the assessment, backfilled ones included; pending backfill patients are left out, as in Zhao et al. (2024). Dose elimination is checked at the same moments, with the same completed patients, at every dose.

Backfilling. A patient who arrives while the escalation waits is backfilled to the highest dose b (or the lowest, with backfill_dose = "lowest") that is below the current dose, has not been

eliminated, has fewer than n_cap patients, and has a response observed at b or at a lower dose. A dose is closed for backfilling when its own DLT rate exceeds the de-escalation boundary and so does the rate pooled with the next higher dose; the doses above it close with it. The rules are evaluated whenever a patient arrives, so a closed dose reopens when new data no longer meet the closing condition.

Conflicts. When the data of a backfilled dose b below the current dose call for a more conservative decision than those of the current dose (Table 2 of Zhao et al., 2024), the decision is taken instead from the data pooled from dose b* to the current dose, where b* is the highest conflicting dose by default. A de-escalation then goes to the highest dose k at or above b* whose data pooled from b* stay at or below the de-escalation boundary, or to b* - 1 if there is none, and to the lowest dose if b* is the lowest.

End of the trial. The trial stops when the escalation has treated n_cohort cohorts, when n_earlystop is reached with a decision to stay, or when the lowest dose is eliminated. Backfilling stops with the escalation. Every enrolled patient is followed to the end of the assessment, and the data of all patients, backfilled ones included, are those from which the MTD is selected.

Arrivals and outcomes. Patients arrive one at a time. The DLT of a patient and its time follow p_true and dlt_time as in tite_boin_simulate; the response and its time follow p_resp, resp_window and resp_late_fraction in the same way, linked to the DLT by resp_cor. Neither article states how responses are timed or correlated with DLTs. The Weibull time to response with half of the responses in the second half of the window and the Gaussian copula follow Takeda et al. (2025), who set the correlation to 0.2; the default here is independence.

Random numbers. The DLT variates of the escalation patients are drawn from R's stream exactly as in boin_simulate; arrivals and the variates of backfill patients and of responses come from two generators of their own seeded with seed. As a result, when no dose is ever opened for backfilling, for example with every p_resp equal to 0, the trials are identical to those of boin_simulate with the same seed.

Agreement with the article. With stay_on_1_of_3 = TRUE, n_earlystop = 9, accrual = "uniform" and no_slot = "leave", the operating characteristics of Table 4 of Zhao et al. (2024) are reproduced within simulation error, durations included. The article states a Poisson process, but its durations, for BOIN as well as for BF-BOIN, agree with uniform times between arrivals and patients turned away while the escalation waits. The trial example of Figure 1 of the article is replayed patient by patient in the tests of the package.

Value

An object of class c("bf_boin_trials", "backfill_trials"), which is a list with components

n_pts	Integer matrix, n_trials by number of doses, of patients treated, backfilled ones included.
n_tox	Integer matrix of the same shape, of DLTs, counted once every patient has completed the assessment.
n_bf	Integer matrix of the same shape, of backfilled patients.
n_resp	Integer matrix of the same shape, of patients with a response within the response window.
eliminated	Logical matrix of the same shape, marking doses eliminated during the trial.
cohorts_used	Integer vector of dose escalation cohorts treated in each trial.

stop_reason	Character vector giving why each trial ended, with the values of boin_simulate.
duration	Numeric vector. Time from the first arrival until every enrolled patient has completed the DLT assessment.
n_suspensions	Integer vector. Number of dose escalation decisions at which the escalation had to wait.
time_suspended	Numeric vector. Total time the escalation waited.
n_turned_away	Integer vector. Patients turned away, which happens only with no_slot = "leave".
boundary	The boin_boundary object used.
settings	List of the design parameters.

References

Zhao, Y., Yuan, Y., Korn, E. L. and Freidlin, B. (2024). Backfilling Patients in Phase I Dose-Escalation Trials Using Bayesian Optimal Interval Design (BOIN). Clinical Cancer Research, 30(4), 673-679.

Chen, K., Zhao, Y., Takeda, K. and Yuan, Y. (2026). BE-BOIN: A Dose Optimization Design Accommodating Backfill and Late-Onset Toxicity. Therapeutic Innovation and Regulatory Science. doi:10.1007/s43441026009940

Takeda, K., Zhu, J. and Hirakawa, A. (2025). BF-BOIN-ET: A Backfill Bayesian Optimal Interval Design Using Efficacy and Toxicity Outcomes for Dose Optimization. Pharmaceutical Statistics, 24(2), e2470.

See Also

be_boin_simulate, boin_simulate

Examples

```
# Scenario 3 of Zhao et al. (2024)
trials <- bf_boin_simulate(
  target = 0.25,
  p_true = c(0.04, 0.12, 0.25, 0.43, 0.63),
  p_resp = c(0.10, 0.20, 0.30, 0.45, 0.58),
  n_cohort = 10,
  cohort_size = 3,
  window = 1,
  accrual_rate = 3,
  n_earlystop = 9,
  stay_on_1_of_3 = TRUE,
  n_trials = 200,
  seed = 123
)
trials

colMeans(trials$n_bf)
```

boin_boundary *Integer Decision Boundaries for the BOIN Design*

Description

Tabulate, for every attainable number of patients treated at a dose, the number of DLTs that triggers escalation, de-escalation, elimination and, when requested, the stricter safety stopping rule at the lowest dose.

Usage

```
boin_boundary(
  target,
  max_n,
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  stay_on_1_of_3 = FALSE
)
```

Arguments

target	Numeric scalar. Target DLT probability, for example 0.30.
max_n	Integer scalar. Largest number of patients that can be treated at a single dose, which is the number of columns of the resulting table. For a trial with n_cohort cohorts of size cohort_size this is n_cohort * cohort_size.
p_saf	Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to 0.6 * target.
p_tox	Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to 1.4 * target.
cutoff_eli	Numeric scalar. Posterior probability cutoff above which a dose is eliminated for toxicity. Defaults to 0.95.
extrasafe	Logical scalar. When TRUE, also tabulate the safety stopping boundary applied at the lowest dose. Defaults to FALSE.
offset	Numeric scalar between 0 and 0.5. Amount by which cutoff_eli is relaxed for the safety stopping rule. Defaults to 0.05.
stay_on_1_of_3	Logical scalar. When TRUE, one DLT out of three patients leads to staying at the current dose rather than de-escalating. Defaults to FALSE. See the details.

Details

The escalation boundary is `floor(lambda_e * n)` and the de-escalation boundary is `ceiling(lambda_d * n)`, with the convention that when `lambda_d * n` is a whole number de-escalation requires one more DLT. The elimination boundary is the smallest positive number of DLTs for which `Pr(p > target | data)` exceeds `cutoff_eli` under a uniform Beta(1, 1) prior, and is not evaluated before three patients have been treated. The de-escalation boundary is capped at the elimination boundary.

With the default thresholds these definitions reproduce `BOIN::get.boundary()` exactly.

Value

An object of class `boin_boundary`, which is a list with components

`lambda_e`	Escalation interval boundary.
`lambda_d`	De-escalation interval boundary.
`n`	Number of patients, 1 to `max_n`.
`b_esc`	Escalate when the number of DLTs is at most this value.
`b_deesc`	De-escalate when the number of DLTs reaches this value.
`b_elim`	Eliminate the dose when the number of DLTs reaches this value, `NA` when no elimination is possible at that sample size.
`b_stop`	Safety stopping boundary at the lowest dose, all `NA` unless `extrasafe` is `TRUE`.
`stay_on_1_of_3`	Whether the modification was requested.
`stay_on_1_of_3_applied`	Whether it changed the boundaries.

together with the design parameters used.

Staying on one DLT out of three

For some target rates the optimal BOIN decision after one DLT in three patients is to de-escalate, because the likelihood of overdosing then exceeds the likelihood of proper dosing. Long practice with the 3+3 design has nonetheless made staying at the dose the widely accepted choice, and `stay_on_1_of_3 = TRUE` aligns the design with that practice by raising the de-escalation boundary at three patients from one DLT to two. Nothing else in the table changes.

The modification is applied only where it is meaningful, that is when one DLT out of three currently triggers de-escalation. It never overrides an escalation, and it never overrides an elimination, which is a safety rule. Whether it took effect is reported in `stay_on_1_of_3_applied`. With the default thresholds it takes effect for target rates from about 0.098 to 0.279; below that range one DLT out of three already eliminates the dose, and above it the design already stays.

References

Liu S. and Yuan, Y. (2015). Bayesian Optimal Interval Designs for Phase I Clinical Trials. Journal of the Royal Statistical Society: Series C, 64, 507-523.

See Also

`boin_lambda`, `boin_decision_table`

Examples

```
bd <- boin_boundary(target = 0.30, max_n = 18)
bd

# At a target of 0.25 one DLT out of three de-escalates by default
boin_boundary(target = 0.25, max_n = 18)$b_deesc[3]

# and stays once the modification is switched on
modified <- boin_boundary(target = 0.25, max_n = 18, stay_on_1_of_3 = TRUE)
modified$b_deesc[3]
modified$stay_on_1_of_3_applied

# At a target of 0.30 the design already stays, so nothing changes
boin_boundary(target = 0.30, max_n = 18, stay_on_1_of_3 = TRUE)$stay_on_1_of_3_applied
```

---

boin_decision_table *Decision Table for the BOIN Design*

Description

Build the lookup table that maps a pair of DLT and patient counts at the current dose to the dose assignment decision.

Usage

```
boin_decision_table(
  target,
  max_n,
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  stay_on_1_of_3 = FALSE
)
```

Arguments

target	Numeric scalar. Target DLT probability.
max_n	Integer scalar. Largest number of patients treated at a single dose.
p_saf	Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to 0.6 * target.
p_tox	Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to 1.4 * target.
cutoff_eli	Numeric scalar. Posterior probability cutoff for dose elimination. Defaults to 0.95.
stay_on_1_of_3	Logical scalar. When TRUE, one DLT out of three patients leads to staying at the current dose rather than de-escalating. Defaults to FALSE. See boin_boundary for what the modification does.

Details

The table is derived from boin_boundary and therefore agrees with the decisions taken by the simulation engine. Elimination takes precedence over the other rules, then escalation, then de-escalation.

Value

An object of class boin_decision_table, which is a character matrix with max_n + 1 rows, labeled by the number of DLTs from 0 to max_n, and max_n columns, labeled by the number of patients from 1 to max_n. Entries are "E" to escalate, "S" to stay, "D" to de-escalate and "DE" to de-escalate and eliminate the dose together with all higher doses. Combinations with more DLTs than patients are NA. The object has print and plot methods; subsetting it returns a plain matrix.

See Also

boin_boundary, plot.boin_decision_table

Examples

```
decisions <- boin_decision_table(target = 0.30, max_n = 12)
decisions

# Only the sample sizes reached at the end of a cohort of three
print(decisions, cohort_size = 3)

# One DLT out of three de-escalates at a target of 0.25
boin_decision_table(target = 0.25, max_n = 9)["1", "3"]

# and stays once the modification is switched on
boin_decision_table(target = 0.25, max_n = 9, stay_on_1_of_3 = TRUE)["1", "3"]
```

---

boin_isotonic *Isotonic Estimates of the DLT Probability*

Description

Estimate the DLT probability at each dose under the constraint that toxicity does not decrease with dose, using the pool adjacent violators algorithm.

Usage

```
boin_isotonic(n_pts, n_tox, admissible = NULL)
```

Arguments

n_pts	Integer matrix with one row per trial and one column per dose, or a vector for a single trial, giving the number of patients treated.
n_tox	Integer matrix or vector of the same shape as n_pts, giving the number of DLTs observed.
admissible	Logical matrix of the same shape as n_pts, or NULL. Doses that are FALSE are left out of the estimation and returned as NA. Defaults to all doses that have treated at least one patient.

Details

Pseudo-counts of 0.05 DLTs and 0.1 patients are added before the fit, and doses are pooled with inverse variance weights. Doses left out through admissible do not influence the estimates of the remaining doses, which matters when eliminated doses must be excluded before selecting the MTD.

Value

A numeric matrix with one row per trial and one column per dose. Doses that did not enter the estimation are NA.

References

Liu S. and Yuan, Y. (2015). Bayesian Optimal Interval Designs for Phase I Clinical Trials. Journal of the Royal Statistical Society: Series C, 64, 507-523.

See Also

boin_select_mtd

Examples

```
# A single trial
boin_isotonic(n_pts = c(3, 6, 9, 12), n_tox = c(0, 1, 3, 4))

# Several trials at once
n_pts <- matrix(c(3, 6, 9, 12,
                  3, 6, 9, 12,
                  3, 6, 9, 12), nrow = 3, byrow = TRUE)
n_tox <- matrix(c(0, 1, 3, 4,
                  0, 0, 2, 3,
                  1, 2, 4, 6), nrow = 3, byrow = TRUE)
boin_isotonic(n_pts, n_tox)
```

`boin_lambda` *Escalation and De-escalation Interval Boundaries*

Description

Compute the two boundaries that define the BOIN decision interval. A dose is escalated when the observed DLT rate falls at or below `lambda_e` and de-escalated when it reaches `lambda_d`.

Usage

```
boin_lambda(target, p_saf = NULL, p_tox = NULL)
```

Arguments

`target`	Numeric scalar. Target DLT probability, for example 0.30.
`p_saf`	Numeric scalar. Highest DLT probability deemed subtherapeutic, so that dose escalation should be undertaken. Defaults to `0.6 * target`.
`p_tox`	Numeric scalar. Lowest DLT probability deemed overly toxic, so that dose de-escalation is required. Defaults to `1.4 * target`.

Details

The boundaries minimize the probability of incorrect dose assignment under a three-point hypothesis on the DLT probability at the current dose. Values of `p_saf` and `p_tox` close to `target` should be avoided, because the sample sizes of phase I trials cannot distinguish the target rate from rates close to it.

Value

A list with components `lambda_e` and `lambda_d`.

References

Liu S. and Yuan, Y. (2015). Bayesian Optimal Interval Designs for Phase I Clinical Trials. Journal of the Royal Statistical Society: Series C, 64, 507-523.

Examples

```
boin_lambda(target = 0.30)

boin_lambda(target = 0.25, p_saf = 0.12, p_tox = 0.40)
```

boin_p_tox *Toxic Threshold Giving a Required De-escalation Boundary*

Description

Find the value of p_tox for which the BOIN de-escalation boundary lambda_d equals a required value. This inverts boin_lambda, which maps p_tox to lambda_d.

Usage

```
boin_p_tox(target, lambda_d)
```

Arguments

target	Numeric scalar. Target DLT probability.
lambda_d	Numeric scalar. Required de-escalation boundary. Must lie strictly between target and 1.

Details

lambda_d increases strictly with p_tox, from target in the limit as p_tox approaches target to 1 in the limit as p_tox approaches 1, so the solution exists and is unique for any lambda_d strictly between those bounds. The root is found numerically.

A boundary close to the target needs a toxic threshold close to the target, and the rest of the package refuses a p_tox within ten percent of target, on the grounds that phase I sample sizes cannot tell the two rates apart. The value is still returned, since the inversion is well defined, but a warning names the smallest boundary that leaves a usable threshold. For a target of 0.30 that smallest boundary is about 0.315.

A design team that wants to tighten the de-escalation boundary usually states the boundary rather than the threshold behind it. This function turns that statement into the p_tox to pass to the other functions, so that the whole design follows from it.

Value

A numeric scalar, the value of p_tox that produces lambda_d.

See Also

boin_lambda, boin_boundary

Examples

```
# A target of 0.30 gives a de-escalation boundary of 0.359 by default
boin_lambda(target = 0.30)$lambda_d

# The toxic threshold that tightens it to 0.33
p_tox <- boin_p_tox(target = 0.30, lambda_d = 0.33)
```

```
p_tox

# which is what it claims to be
boin_lambda(target = 0.30, p_tox = p_tox)$lambda_d

# A boundary too close to the target leaves an unusable threshold, and says so
boin_p_tox(target = 0.30, lambda_d = 0.31)

# and shifts the de-escalation boundary down by one DLT at every cohort end
as.data.frame(boin_boundary(0.30, 18))[seq(3, 18, 3), ]
as.data.frame(boin_boundary(0.30, 18, p_tox = p_tox))[seq(3, 18, 3), ]
```

`boin_select_mtd` *Select the MTD from Completed Trials*

Description

Choose the maximum tolerated dose for each of a set of completed trials from the final patient and DLT counts.

Usage

```
boin_select_mtd(
  n_pts,
  n_tox,
  target,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  bound_mtd = FALSE,
  p_tox = NULL,
  mtd_max_estimate = NULL,
  min_mtd_sample = 1
)
```

Arguments

`n_pts`	Integer matrix with one row per trial and one column per dose, or a vector for a single trial, giving the number of patients treated.
`n_tox`	Integer matrix or vector of the same shape as `n_pts`, giving the number of DLTs observed.
`target`	Numeric scalar. Target DLT probability.
`cutoff_eli`	Numeric scalar. Posterior probability cutoff for dose elimination. Defaults to 0.95.
`extrasafe`	Logical scalar. Apply the stricter safety rule at the lowest dose, in which case no dose is selected when the lowest dose is judged too toxic. Defaults to `FALSE`.

`offset` Numeric scalar between 0 and 0.5. Amount by which `cutoff_eli` is relaxed for the safety rule. Defaults to 0.05.

`bound_mtd` Logical scalar. Require the isotonic estimate at the selected dose to be at or below the de-escalation boundary. Defaults to `FALSE`.

`p_tox` Numeric scalar. Lowest DLT probability deemed overly toxic, used to derive the de-escalation boundary when `bound_mtd` is `TRUE`. Defaults to `1.4 * target`.

`mtd_max_estimate` Numeric scalar or `NULL`. Largest isotonic estimate a dose may have and still be selected. Supplying it bounds the selection whatever `bound_mtd` says, and replaces the de-escalation boundary that `bound_mtd` would otherwise use. Defaults to `NULL`.

`min_mtd_sample` Integer scalar. Smallest number of patients a dose must have received to be eligible. Defaults to 1, which admits every dose that treated a patient.

Details

Dose elimination is re-derived from the final data of each trial rather than carried over from the dose-finding stage, which is what the reference implementation does and which matters when a trial ends at its maximum sample size. The isotonic fit is computed over the admissible doses only, so that eliminated doses do not influence the estimates of the doses that remain. The dose whose estimate is closest to `target` is selected, with ties broken by a small increasing perturbation across the admissible doses.

`bound_mtd` caps the estimate at the de-escalation boundary, which always lies above `target`. A cap at or below the target rate therefore cannot be expressed that way, and needs `mtd_max_estimate`. The two caps are independent of the decision table, so bounding the selection does not change how the trial was run.

Values of `reason` are `"selected"`, `"lowest_dose_eliminated"`, `"no_admissible_dose"` and `"no_dose_below_bound"`.

Value

A data frame with one row per trial and columns

`trial` Trial index.

`mtd` Selected dose level, or `NA` when no dose is selected.

`reason` Why the dose was or was not selected.

References

Liu S. and Yuan, Y. (2015). Bayesian Optimal Interval Designs for Phase I Clinical Trials. Journal of the Royal Statistical Society: Series C, 64, 507-523.

See Also

`boin_isotonic`, `boin_simulate`

Examples

```
n_pts <- matrix(c(3, 6, 9, 3,
                  3, 6, 9, 3), nrow = 2, byrow = TRUE)
n_tox <- matrix(c(0, 1, 3, 2,
                  0, 1, 2, 1), nrow = 2, byrow = TRUE)

boin_select_mtd(n_pts, n_tox, target = 0.30)

boin_select_mtd(n_pts, n_tox, target = 0.30, bound_mtd = TRUE)

# Require the isotonic estimate at the selected dose to be at most 0.30,
# which the de-escalation boundary cannot express
boin_select_mtd(n_pts, n_tox, target = 0.30, mtd_max_estimate = 0.30)
```

`boin_simulate` *Simulate BOIN Trials*

Description

Run the BOIN dose-finding algorithm over many simulated trials and return the patient and DLT counts at every dose. This is the simulation engine behind `sim_boin`, exposed for users who want the raw trial data.

Usage

```
boin_simulate(
  target,
  p_true,
  n_cohort,
  cohort_size,
  n_trials = 10000,
  start_dose = 1,
  n_earlystop = 18,
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  titration = FALSE,
  stay_on_1_of_3 = FALSE,
  n_earlystop_rule = c("with_stay", "simple"),
  seed = 123
)
```

Arguments

target	Numeric scalar. Target DLT probability, for example 0.30.
p_true	Numeric vector. True DLT probability at each dose level, in increasing dose order.
n_cohort	Integer scalar. Number of cohorts in a trial.
cohort_size	Integer scalar or vector. Number of patients per cohort. A scalar is used for every cohort. A vector shorter than n_cohort is padded with its last element and a longer one is truncated.
n_trials	Integer scalar. Number of trials to simulate. Defaults to 10000.
start_dose	Integer scalar. Dose level for the first cohort. Defaults to 1. It is ignored when titration is TRUE, because the titration phase always begins at the lowest dose.
n_earlystop	Integer scalar. The trial stops once this many patients have been treated at the current dose and the design would stay there. Defaults to 18.
p_saf	Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to 0.6 * target.
p_tox	Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to 1.4 * target.
cutoff_eli	Numeric scalar. Posterior probability cutoff for dose elimination. Defaults to 0.95.
extrasafe	Logical scalar. Apply the stricter safety stopping rule at the lowest dose. Defaults to FALSE.
offset	Numeric scalar between 0 and 0.5. Amount by which cutoff_eli is relaxed for the safety stopping rule. Defaults to 0.05.
titration	Logical scalar. Start with single patient cohorts until the first DLT is seen. Ignored when the first cohort size is one. Defaults to FALSE.
stay_on_1_of_3	Logical scalar. When TRUE, one DLT out of three patients leads to staying at the current dose rather than de-escalating. Defaults to FALSE. See boin_boundary.
n_earlystop_rule	Character scalar, either "with_stay" or "simple". Under "with_stay" the trial stops at n_earlystop only when the design would remain at the current dose, which is the rule used by the reference implementation. Under "simple" it stops as soon as the threshold is reached.
seed	Integer scalar or NULL. Random seed. When NULL the current state of the random number generator is used. The state of the calling session is restored on exit. Defaults to 123.

Details

One uniform variate is drawn per patient, in enrollment order. Two details matter for reproducibility: the titration phase draws one variate per dose level whether or not it is used, and a cohort cut short by the maximum sample size still draws a full cohort's worth. This is the random number consumption of BOIN::get.oc(), so with the same seed and matching arguments a run reproduces the reference implementation trial by trial.

Possible values of stop_reason are "lowest_dose_eliminated", "lowest_dose_too_toxic", "n_earlystop", "max_sample_size" and "max_cohorts".

Value

An object of class `boin_trials`, which is a list with components

`n_pts`	Integer matrix, `n_trials` by number of doses, of patients treated.
`n_tox`	Integer matrix of the same shape, of DLTs observed.
`eliminated`	Logical matrix of the same shape, marking doses eliminated during the trial.
`cohorts_used`	Integer vector of cohorts completed in each trial.
`stop_reason`	Character vector giving why each trial ended.
`boundary`	The `boin_boundary` object used.
`settings`	List of the design parameters.

References

Liu S. and Yuan, Y. (2015). Bayesian Optimal Interval Designs for Phase I Clinical Trials. Journal of the Royal Statistical Society: Series C, 64, 507-523.

See Also

`sim_boin`, `boin_select_mtd`

Examples

```
trials <- boin_simulate(
  target = 0.30,
  p_true = c(0.05, 0.15, 0.30, 0.45, 0.60),
  n_cohort = 10,
  cohort_size = 3,
  n_trials = 200,
  seed = 123
)
trials

head(trials$n_pts)
table(trials$stop_reason)
```

---

boin_stopping_table *Safety Stopping Boundary as a Two-Row Table*

Description

Present the safety stopping boundary at the lowest dose as a compact table with one column per sample size, which is the shape usually wanted in a protocol or a report.

Usage

```
boin_stopping_table(x, cohort_size = NULL)
```

Arguments

x	An object of class boin_boundary, built with extrasafe = TRUE.
cohort_size	Integer scalar or NULL. When supplied, only the sample sizes that are multiples of cohort_size are kept.

Details

The same numbers are available one row per sample size from as.data.frame() on the boundary object. This function transposes them and labels the rows, which suits a table pasted into a document or rendered with knitr::kable() or DT::datatable().

Value

A data frame with two rows, labeled by the quantity they hold, and one column per sample size. Sample sizes at which no number of DLTs triggers the rule are dropped.

See Also

boin_boundary

Examples

```
bd <- boin_boundary(target = 0.30, max_n = 18, extrasafe = TRUE)

boin_stopping_table(bd)

# Only the sample sizes reached at the end of a cohort of three
boin_stopping_table(bd, cohort_size = 3)
```

---

comb_boin_next *Next Dose Combination in a BOIN Drug-Combination Trial*

Description

Decide the dose combination for the next cohort of a trial run with the BOIN drug-combination design of Lin and Yin (2017), from the numbers of patients and DLTs observed so far at every combination and the combination that treated the last cohort. The rules are those of comb_boin_simulate, so the decisions taken during the trial are the ones whose operating characteristics sim_comb_boin reports.

Usage

```
comb_boin_next(
  target,
  n_pts,
  n_tox,
```

```
  current,
  n_earlystop = 18,
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05
)
```

Arguments

target	Numeric scalar. Target DLT probability, for example 0.30.
n_pts	Integer matrix. Number of patients treated so far at each dose combination, with one row per dose level of drug A and one column per dose level of drug B.
n_tox	Integer matrix of the same shape as n_pts. Number of patients with a DLT at each combination.
current	Integer vector of length two. Dose levels of drug A and drug B of the combination that treated the last cohort.
n_earlystop	Integer scalar. The trial stops once this many patients have been treated at the current combination and the design would neither escalate nor de-escalate, or cannot. Defaults to 18. Use the value of the simulations that support the protocol.
p_saf	Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to 0.6 * target.
p_tox	Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to 1.4 * target.
cutoff_eli	Numeric scalar. Posterior probability cutoff for eliminating a combination. Defaults to 0.95.
extrasafe	Logical scalar. Apply the stricter safety stopping rule at the lowest combination. Defaults to FALSE.
offset	Numeric scalar between 0 and 0.5. Amount by which cutoff_eli is relaxed for the safety stopping rule. Defaults to 0.05.

Details

The decision follows the same steps as a simulated trial after each cohort. A combination whose DLT count reaches its elimination boundary is eliminated together with every combination with at least the same level of both drugs, and the trial stops if the lowest combination is eliminated or, with extrasafe, judged too toxic. The trial also stops when n_earlystop patients have been treated at the current combination and the design would neither escalate nor de-escalate. Otherwise the observed DLT rate at the current combination is compared with the boundaries of boin_boundary, and an escalation or a de-escalation moves one level of drug A or of drug B to the candidate with the higher score: the posterior probability, under a Beta(0.5, 0.5) prior, that its DLT rate lies between lambda_e and lambda_d, plus 0.0005 per patient already treated there. An eliminated combination is not a candidate for escalation. When the two scores are equal the design chooses at random, and the function returns both.

During a simulated trial a combination is eliminated as soon as its data reach the boundary and is never treated again, so the combinations eliminated by the data are exactly those eliminated in the course of the trial.

`BOIN::next.comb()` applies somewhat different rules from the simulations of `BOIN::get.oc.comb()`: it stops as soon as `n.earlystop` patients have been treated at the current combination, it does not escalate drug A when a combination with the next level of drug A and at most the current level of drug B has an observed DLT rate at or above `lambda_d` (and likewise for drug B), and it draws a random number to break ties. This function follows the simulations instead, so that a trial is run with the rules whose operating characteristics were evaluated. Where the rules coincide, the two functions agree.

Value

An object of class `comb_boin_next`, which is a list with components

`decision`	One of `"escalate"`, `"de-escalate"`, `"stay"` and `"stop"`.
`next_dose`	Integer matrix with columns `dose_a` and `dose_b`: one row for the next combination, two rows when two candidates tie, and no row when the trial stops.
`tie`	Whether two candidates tie, in which case the design picks one of them at random.
`stop_reason`	`NA`, or why the trial stops: `"lowest_dose_eliminated"`, `"lowest_dose_too_toxic"` or `"n_earlystop"`.
`note`	A sentence explaining a decision to stay when the rule called for a move that is not available, otherwise `NA`.
`candidates`	Data frame of the two combinations considered for an escalation or a de-escalation, one level of drug A and one level of drug B away, with their data, the posterior probability that their DLT rate lies between `lambda_e` and `lambda_d`, the score used to choose between them and whether they are available; `NULL` when no move was considered.
`eliminated`	Logical matrix of the combinations eliminated by the data.
`boundary`	The escalation, de-escalation and elimination boundaries at the number of patients treated at the current combination.

together with the data, the current combination and the design parameters.

References

Lin, R. and Yin, G. (2017). Bayesian Optimal Interval Design for Dose Finding in Drug-Combination Trials. Statistical Methods in Medical Research, 26(5), 2155-2167.

See Also

`plot.comb_boin_next`, `boin_decision_table`, `sim_comb_boin`

Examples

```
# Three levels of drug A and four of drug B; the last cohort was treated at
# A2B2, where two of three patients had a DLT
n_pts <- matrix(c(6, 6, 0, 0,
                  3, 3, 0, 0,
                  0, 0, 0, 0), nrow = 3, byrow = TRUE)
n_tox <- matrix(c(0, 1, 0, 0,
                  1, 2, 0, 0,
                  0, 0, 0, 0), nrow = 3, byrow = TRUE)

nxt <- comb_boin_next(target = 0.25, n_pts = n_pts, n_tox = n_tox,
                      current = c(2, 2))
nxt
nxt$next_dose
```

---

comb_boin_select_mtd *Select the MTD Combination from Completed Drug-Combination Trials*

Description

Choose the maximum tolerated dose combination for each of a set of completed trials of the BOIN drug-combination design from the final patient and DLT counts.

Usage

```
comb_boin_select_mtd(
  n_pts,
  n_tox,
  target,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  bound_mtd = FALSE,
  p_tox = NULL
)
```

Arguments

n_pts	Integer matrix giving the number of patients treated at each combination of one trial, with one row per dose level of drug A and one column per dose level of drug B, or an integer array with dimensions trial by level of drug A by level of drug B for several trials, as returned by comb_boin_simulate.
n_tox	Integer matrix or array of the same shape as n_pts, giving the number of DLTs observed.
target	Numeric scalar. Target DLT probability.

cutoff_eli	Numeric scalar. Posterior probability cutoff for eliminating a combination. Defaults to 0.95.
extrasafe	Logical scalar. Apply the stricter safety rule at the lowest combination, in which case no combination is selected when it is judged too toxic. Defaults to FALSE.
offset	Numeric scalar between 0 and 0.5. Amount by which cutoff_eli is relaxed for the safety rule. Defaults to 0.05.
bound_mtd	Logical scalar. Require the isotonic estimate at the selected combination to be below the de-escalation boundary. Defaults to FALSE.
p_tox	Numeric scalar. Lowest DLT probability deemed overly toxic, used to derive the de-escalation boundary when bound_mtd is TRUE. Defaults to 1.4 * target.

Details

The selection follows the reference implementation used by BOIN::get.oc.comb(). Elimination is re-derived from the final data: within each row, the first combination with at least three patients whose posterior probability of a DLT rate above target exceeds cutoff_eli is eliminated together with the rest of its row and the rest of its column. If the lowest combination is eliminated, no combination is selected. Otherwise the estimates (y + 0.05) / (n + 0.1) are smoothed by bivariate isotonic regression with weights n + 0.1, so that they increase along every row and column, and the treated, non-eliminated combination whose estimate is closest to target is selected. A small amount, 1e-5 times the sum of the two dose levels, is added to every estimate first, so that ties go to the lower combination.

The isotonic regression is Iso::biviso(), the routine the reference implementation calls. Combinations that the regression pools together have equal estimates in exact arithmetic but can differ in the last digits, and with the small amount above, these digits decide between combinations on the same anti-diagonal. Another algorithm for the same regression would therefore select differently in a few trials out of a thousand.

The article of Lin and Yin (2017) estimates the DLT rates as y / n and resolves ties differently; the package follows the reference implementation so that the two can be compared trial by trial.

Values of reason are "selected", "lowest_dose_eliminated" and "no_dose_below_bound".

Value

A data frame with one row per trial and columns

trial	Trial index.
mtd_a	Dose level of drug A in the selected combination, or NA when none is selected.
mtd_b	Dose level of drug B in the selected combination, or NA.
reason	Why a combination was or was not selected.

References

Lin, R. and Yin, G. (2017). Bayesian Optimal Interval Design for Dose Finding in Drug-Combination Trials. Statistical Methods in Medical Research, 26(5), 2155-2167.

See Also

comb_boin_simulate, sim_comb_boin

Examples

```
n_pts <- matrix(c(3, 6, 3,
                  3, 3, 0), nrow = 2, byrow = TRUE)
n_tox <- matrix(c(0, 1, 1,
                  0, 2, 0), nrow = 2, byrow = TRUE)

comb_boin_select_mtd(n_pts, n_tox, target = 0.30)
```

---

comb_boin_simulate *Simulate BOIN Drug-Combination Trials*

Description

Run the BOIN drug-combination design of Lin and Yin (2017) over many simulated trials of two drugs given together, and return the patient and DLT counts at every dose combination. This is the simulation engine behind sim_comb_boin, exposed for users who want the raw trial data.

Usage

```
comb_boin_simulate(
  target,
  p_true,
  n_cohort,
  cohort_size,
  n_trials = 10000,
  start_dose = c(1, 1),
  n_earlystop = 18,
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  titration = FALSE,
  seed = 123
)
```

Arguments

target	Numeric scalar. Target DLT probability, for example 0.30.
p_true	Numeric matrix. True DLT probability of each dose combination, with one row per dose level of drug A and one column per dose level of drug B, both in increasing dose order.

n_cohort	Integer scalar. Number of cohorts in a trial.
cohort_size	Integer scalar or vector. Number of patients per cohort. A scalar is used for every cohort. A vector shorter than n_cohort is padded with its last element and a longer one is truncated.
n_trials	Integer scalar. Number of trials to simulate. Defaults to 10000.
start_dose	Integer vector of length two. Dose levels of drug A and drug B for the first cohort. Defaults to c(1, 1). Titration, when requested, also starts here.
n_earlystop	Integer scalar. The trial stops once this many patients have been treated at the current combination and the design would neither escalate nor de-escalate, or cannot. Defaults to 18. Set it to a value above the maximum sample size to switch this rule off.
p_saf	Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to 0.6 * target.
p_tox	Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to 1.4 * target.
cutoff_eli	Numeric scalar. Posterior probability cutoff for eliminating a combination. Defaults to 0.95.
extrasafe	Logical scalar. Apply the stricter safety stopping rule at the lowest combination. Defaults to FALSE.
offset	Numeric scalar between 0 and 0.5. Amount by which cutoff_eli is relaxed for the safety stopping rule. Defaults to 0.05.
titration	Logical scalar. Treat one patient per combination, moving up one level of either drug, until the first DLT, before switching to full cohorts. Ignored when the first cohort size is one. Defaults to FALSE.
seed	Integer scalar or NULL. Random seed. When NULL the current state of the random number generator is used. The state of the calling session is restored on exit. Defaults to 123.

Details

The decision boundaries are those of the single-agent design, from boin_boundary. After each cohort the observed DLT rate at the current combination is compared with them. When the design escalates, the next cohort goes one level up of drug A or of drug B, whichever has the larger posterior probability, under a Beta(0.5, 0.5) prior, that its DLT probability lies between lambda_e and lambda_d; de-escalation chooses one level down in the same way. Eliminated combinations are not escalated to. When a combination is eliminated, so is every combination with at least the same level of both drugs, and the trial stops when the lowest combination is eliminated.

The rules, and the order in which they are applied, are those of BOIN::get.oc.comb() with mtd.contour = FALSE. That implementation adds two things to the article: a candidate's posterior probability is increased by 0.0005 for every patient already treated there, so that ties go to the better informed combination, and the trial can stop at n_earlystop. Random variates are drawn in the same order as there: one per patient, one for the direction of every titration step that can go either way, and one for every escalation or de-escalation that has a candidate, even when there is only one. With the same seed and matching arguments a run therefore reproduces the reference implementation trial by trial.

During titration the patients treated before the first DLT are added to the planned cohorts, as in the reference implementation: the first full cohort only tops up the combination reached. A trial can therefore treat a few more patients than `n_cohort` cohorts hold.

Possible values of `stop_reason` are `"lowest_dose_eliminated"`, `"lowest_dose_too_toxic"`, `"n_earlystop"` and `"max_cohorts"`.

Value

An object of class `comb_boin_trials`, which is a list with components

`n_pts`	Integer array of patients treated, with dimensions trial by level of drug A by level of drug B.
`n_tox`	Integer array of the same shape, of DLTs observed.
`eliminated`	Logical array of the same shape, marking combinations eliminated during the trial.
`cohorts_used`	Integer vector of cohorts completed in each trial.
`stop_reason`	Character vector giving why each trial ended.
`boundary`	The `boin_boundary` object used.
`settings`	List of the design parameters.

References

Lin, R. and Yin, G. (2017). Bayesian Optimal Interval Design for Dose Finding in Drug-Combination Trials. Statistical Methods in Medical Research, 26(5), 2155-2167.

Yan, F., Zhang, L., Zhou, Y., Pan, H., Liu, S. and Yuan, Y. (2020). BOIN: An R Package for Designing Single-Agent and Drug-Combination Dose-Finding Trials Using Bayesian Optimal Interval Designs. Journal of Statistical Software, 94(13), 1-32.

See Also

`sim_comb_boin`, `comb_boin_select_mtd`

Examples

```
p_true <- matrix(c(0.02, 0.04, 0.08, 0.14,
                   0.08, 0.25, 0.42, 0.48,
                   0.25, 0.45, 0.50, 0.60), nrow = 3, byrow = TRUE)

trials <- comb_boin_simulate(
  target = 0.25,
  p_true = p_true,
  n_cohort = 16,
  cohort_size = 3,
  n_trials = 200,
  seed = 123
)
trials
```

```
# Patients treated at each combination in the first trial
trials$n_pts[1, , ]
```

`decision_table_3p3` *Decision Table for the 3+3 Design*

Description

Tabulate the dose assignment rule of the traditional 3+3 design used by `oc_3p3` and `sim_3p3`: for every number of DLTs among the three or six patients treated at the current dose, the decision taken during dose escalation and, for `mtd_rule = "expand"`, during the search for the MTD that follows it.

Usage

```
decision_table_3p3(mtd_rule = c("previous", "expand"))
```

Arguments

`mtd_rule`	Character scalar, either `"previous"` or `"expand"`. The definition of the MTD, as in `oc_3p3`. Defaults to `"previous"`.

Details

During escalation three patients are treated at the current dose. No DLT escalates, one DLT adds three more patients at the same dose, and two or more stop escalation. With six patients, at most one DLT escalates and two or more stop escalation. Escalating from the highest dose also ends escalation.

With `mtd_rule = "previous"` the MTD is the dose below the one at which escalation stopped, or the highest dose when it was escalated from, and no MTD is selected when escalation stops at the starting dose. With `mtd_rule = "expand"` the search starts at the dose below the one at which escalation stopped, or at the highest dose when it was escalated from. A dose with three patients receives three more; a dose with at most one DLT out of six is selected as the MTD, and one with two or more passes the search to the next lower dose. No MTD is selected when the search moves below the starting dose.

Every number of DLTs from 0 to `n` is listed for six patients, although in a trial the six patients at a dose during escalation always include one DLT among the first three. The search examines only doses that were escalated from, so a dose with three patients has no DLT there, and only that state is listed.

The rules are those of `oc_3p3`: the tests follow this table through every possible course of a trial and obtain the operating characteristics of `oc_3p3()`.

Value

An object of class `decision_table_3p3`, which is a data frame with one row per state of the current dose and the columns

`stage`	`"escalation"` for dose escalation, or `"search"` for the search for the MTD, which only `mtd_rule = "expand"` has.
`n`	Number of patients treated at the current dose, 3 or 6.
`n_tox`	Number of patients with a DLT, from 0 to `n`.
`decision`	Decision code. During escalation, `"E"` escalates to the next higher dose, `"S"` treats three more patients at the current dose and `"STOP"` stops escalation, this dose and all higher doses being too toxic. During the search, `"S"` treats three more patients at the current dose, `"MTD"` selects it as the MTD and `"D"` moves the search to the next lower dose.

The definition of the MTD is stored in the attribute `mtd_rule`. The object has `print` and `plot` methods.

See Also

`print.decision_table_3p3`, `plot.decision_table_3p3`, `oc_3p3`, `boin_decision_table`

Examples

```
decision_table_3p3()

# The other definition of the MTD adds the search for the MTD
decision_table_3p3(mtd_rule = "expand")
```

---

oc_3p3 *Exact Operating Characteristics of the 3+3 Design*

Description

Compute the operating characteristics of the traditional 3+3 dose-finding design exactly. The probability of every path the design can take is evaluated in closed form, so the result carries no Monte Carlo error and no seed is involved. Use it as a comparator for the BOIN results returned by `sim_boin`.

Usage

```
oc_3p3(
  p_true,
  mtd_rule = c("previous", "expand"),
  start_dose = 1,
  overdose_cutoff = NULL
)
```

Arguments

p_true	Numeric vector. True DLT probability at each dose level, in increasing dose order.
mtd_rule	Character scalar, either "previous" or "expand". The two definitions of the MTD in common use. See Details. Defaults to "previous".
start_dose	Integer scalar. Dose level for the first cohort. Defaults to 1.
overdose_cutoff	Numeric scalar or NULL. Doses whose true DLT probability exceeds this value count as overdoses in the overdose component of the result. Defaults to NULL, which uses one third.

Details

The escalation rule is the traditional one. Three patients are treated at the current dose. No DLT leads to the next dose up. One DLT out of three leads to three more patients at the same dose, and then to the next dose up if no further DLT is seen and to stopping otherwise. Two or more DLTs out of three stop the trial at once. A trial that escalates past the highest dose ends there.

The two MTD rules differ only in what happens once escalation has stopped. "previous" names the dose below the one declared too toxic, and the highest dose when escalation ran out of doses. "expand" instead takes the MTD to be the highest dose at which at most one of six patients had a DLT: starting from the highest dose still eligible, a dose that has only three patients is expanded to six and assessed, and the search moves down a level whenever the expanded dose fails. Neither rule is universally the standard, so the comparator being matched should be checked before a rule is chosen. The two agree on which doses are visited during escalation, and so on the probability that a patient is exposed to an overly toxic dose; they differ in the sample size and in the selection percentages.

The 3+3 design has no fixed maximum sample size, so the within-trial exposure percentage and the 60 percent and 80 percent thresholds reported for BOIN in sim_boin have no counterpart here.

Value

An object of class oc_3p3, which is a list with components

sel_percent	Percentage of trials selecting each dose as the MTD.
percent_no_mtd	Percentage of trials ending without an MTD.
n_pts_dose	Average number of patients treated at each dose.
n_tox_dose	Average number of DLTs observed at each dose.
total_n_pts	Average total number of patients per trial.
total_n_tox	Average total number of DLTs per trial.
overdose	List describing the design's relationship with doses above overdose_cutoff: pct_patients (the percentage of all patients treated at such a dose, that is the probability that a patient is dosed above the cutoff), avg_n_patients, pct_trials_any (the percentage of trials treating at least one patient there), pct_trials_mtd_above (the percentage of all trials whose selected MTD is above the cutoff) and pct_mtd_above_when_selected (the same among the trials that selected an MTD). The cutoff and the dose levels doses that exceed it are also returned.

together with p_true, n_doses, mtd_rule, start_dose, method and the call.

See Also

sim_3p3, sim_boin

Examples

```
oc_3p3(p_true = c(0.30, 0.48, 0.67))

# The other definition of the MTD
oc_3p3(p_true = c(0.30, 0.48, 0.67), mtd_rule = "expand")

# Exposure to doses above a stated DLT rate
oc_3p3(p_true = c(0.20, 0.40, 0.60), overdose_cutoff = 1 / 3)$overdose
```

---

plot.boin_decision_table

Plot a BOIN Decision Table

Description

Draw the decision table produced by boin_decision_table as a grid of colored cells, one per pair of DLT and patient counts, with the decision letter in each cell.

Usage

```
## S3 method for class 'boin_decision_table'
plot(x, text_size = 1, colors = NULL, ...)
```

Arguments

x	An object of class boin_decision_table.
text_size	Numeric scalar. Multiplier applied to every text size in the figure. Use a larger value when the figure is exported at high resolution, for example into a report. Defaults to 1.
colors	Named character vector of four colors, with names "E", "S", "D" and "DE". Defaults to a palette whose adjacent categories remain distinguishable under the common forms of color vision deficiency.
...	Further arguments, currently ignored.

Details

Every cell carries its decision letter, so the figure does not rely on color alone to convey the decision and remains readable in grayscale.

Requires the **ggplot2** package, which is only suggested by **simFastBOIN** rather than required.

Value

A **ggplot2** object, which can be printed, saved with `ggplot2::ggsave()` or extended with further layers.

See Also

`boin_decision_table`, `print.boin_decision_table`

Examples

```
decisions <- boin_decision_table(target = 0.30, max_n = 12)

plot(decisions)

# Larger text for a high resolution export
plot(decisions, text_size = 2)

# A palette of your own
plot(decisions, colors = c(E = "#4DAF4A", S = "#377EB8",
                           D = "#FF7F00", DE = "#E41A1C"))
```

---

plot.comb_boin_next *Plot the Next Dose Combination of a BOIN Drug-Combination Trial*

Description

Draw the dose matrix of a drug-combination trial with the data observed at every combination, the current combination, the combination chosen for the next cohort and the combinations eliminated, as returned by `comb_boin_next`.

Usage

```
## S3 method for class 'comb_boin_next'
plot(x, text_size = 1, colors = NULL, ...)
```

Arguments

`x`	An object of class `comb_boin_next`.
`text_size`	Numeric scalar. Multiplier applied to every text size in the figure. Use a larger value when the figure is exported at high resolution. Defaults to 1.
`colors`	Named character vector of five colors, with names `"Next"`, `"Current"`, `"Eliminated"`, `"Treated"` and `"Not treated"`. Defaults to the palette of `plot.boin_decision_table`, with gray and white for the combinations that are neither current, next nor eliminated.
`...`	Further arguments, currently ignored.

Details

Each cell shows the number of patients with a DLT and the number treated, and the status of the current, next and eliminated combinations is also written in the cell, so the figure does not rely on color alone. The candidates of an escalation or a de-escalation also show the posterior probability that their DLT rate lies between `lambda_e` and `lambda_d`. The levels of drug A run up the vertical axis and those of drug B along the horizontal axis.

Requires the **ggplot2** package, which is only suggested by **simFastBOIN** rather than required.

Value

A **ggplot2** object, which can be printed, saved with `ggplot2::ggsave()` or extended with further layers.

See Also

`comb_boin_next`

Examples

```
n_pts <- matrix(c(6, 6, 0, 0,
                  3, 3, 0, 0,
                  0, 0, 0, 0), nrow = 3, byrow = TRUE)
n_tox <- matrix(c(0, 1, 0, 0,
                  1, 2, 0, 0,
                  0, 0, 0, 0), nrow = 3, byrow = TRUE)

nxt <- comb_boin_next(target = 0.25, n_pts = n_pts, n_tox = n_tox,
                      current = c(2, 2))
plot(nxt)
```

---

plot.decision_table_3p3

Plot a 3+3 Decision Table

Description

Draw the decision table produced by `decision_table_3p3` as a grid of colored cells, one per pair of patient and DLT counts, with the decision code in each cell. With `mtd_rule = "expand"` the search for the MTD is drawn in a second panel.

Usage

```
## S3 method for class 'decision_table_3p3'
plot(x, text_size = 1, colors = NULL, ...)
```

Arguments

x	An object of class decision_table_3p3.
text_size	Numeric scalar. Multiplier applied to every text size in the figure. Defaults to 1.
colors	Named character vector of colors with the names "E", "S" and "STOP", and also "MTD" and "D" for mtd_rule = "expand". Defaults to the palette of plot.boin_decision_table, with its color for elimination used for "STOP" and a violet for "MTD".
...	Further arguments, currently ignored.

Details

Every cell carries its decision code, so the figure does not rely on color alone. Requires the **ggplot2** package, which is only suggested by **simFastBOIN** rather than required.

Value

A **ggplot2** object, which can be printed, saved with ggplot2::ggsave() or extended with further layers.

See Also

decision_table_3p3, print.decision_table_3p3

Examples

```
plot(decision_table_3p3())

plot(decision_table_3p3(mtd_rule = "expand"))
```

---

plot.tite_boin_decision_table

Plot a TITE-BOIN Decision Table

Description

Draw the decision table produced by tite_boin_decision_table as one panel per number of patients treated. Within a panel, each cell is a pair of pending patients and DLTs, colored by its decision and labeled with the decision code and, where the decision depends on the follow-up of the pending patients, the boundaries on the follow-up statistic.

Usage

```
## S3 method for class 'tite_boin_decision_table'
plot(
  x,
  n = NULL,
  cohort_size = NULL,
  text_size = 1,
  colors = NULL,
  digits = 2,
  ...
)
```

Arguments

x	An object of class tite_boin_decision_table.
n	Integer vector or NULL. Numbers of patients treated to draw, one panel each. Defaults to every number in the table.
cohort_size	Integer scalar or NULL. When supplied, only the numbers of patients that are multiples of cohort_size are drawn.
text_size	Numeric scalar. Multiplier applied to every text size in the figure. Defaults to 1.
colors	Named character vector of six colors, with names "E", "S", "D", "DE", "SUS" and "Depends", the last for the states whose decision depends on the follow-up statistic. Defaults to the palette of plot.boin_decision_table extended with gray for suspension and a pale violet for the follow-up dependent states.
digits	Integer scalar. Number of decimal places of the boundaries in the labels. Defaults to 2.
...	Further arguments, currently ignored.

Details

A cell whose decision depends on the follow-up statistic shows its codes over the boundaries, kept short so that they fit in the cell. A label such as "E/S" over ">=2.15" means that the dose is escalated when the statistic is at least 2.15 and kept otherwise, and "S/D" over "<=0.52" that it is de-escalated when the statistic is at most 0.52. The boundary marked ">=" always belongs to the first code, which is "E", or "SUS" when escalation is blocked, and the one marked "<=" to "D". The statistic is STFT for method = "imputation" and ESS for method = "ess".

When the table was built with a positive min_follow_up, the code of a cell whose escalation also requires the shortest follow-up of the pending patients to reach that fraction of the window is marked with an asterisk, as in "E*" or "E/S*", and a caption explains the mark.

Every cell carries its decision code, so the figure does not rely on color alone. With many panels, or many patients per panel, the cells become small; draw fewer panels through n, enlarge the figure, or lower text_size. Requires the **ggplot2** package, which is only suggested by **simFastBOIN** rather than required.

Value

A **ggplot2** object, which can be printed, saved with `ggplot2::ggsave()` or extended with further layers.

See Also

`tite_boin_decision_table`, `print.tite_boin_decision_table`

Examples

```
decisions <- tite_boin_decision_table(target = 0.2, max_n = 15)

plot(decisions, n = c(9, 12, 15))

# Every cohort end, with smaller text
plot(decisions, cohort_size = 3, text_size = 0.8)
```

---

print.backfill_oc *Print Operating Characteristics of a BOIN Design with Backfilling*

Description

Display the result of `sim_bf_boin` or `sim_be_boin`: the design, a table by dose of the true DLT and response rates, the MTD selection, the patients treated and backfilled and the DLTs, then the exposure above the overdose cutoff and the timing.

Usage

```
## S3 method for class 'backfill_oc'
print(
  x,
  digits = 1,
  percent = FALSE,
  kable = FALSE,
  kable_format = "pipe",
  ...
)
```

Arguments

`x`	An object of class `backfill_oc`.
`digits`	Integer scalar. Number of decimal places. Defaults to 1.
`percent`	Logical scalar. Show patients, backfilled patients and DLTs as percentages of their totals rather than as averages per trial. Defaults to `FALSE`.
`kable`	Logical scalar. Print only the table, formatted by `knitr::kable()`. Defaults to `FALSE`.

kable_format	Character scalar passed to knitr::kable() as format. Defaults to "pipe".
...	Further arguments, currently ignored.

Value

The object x, invisibly.

See Also

sim_bf_boin, sim_be_boin

Examples

```
oc <- sim_bf_boin(
  target = 0.25,
  p_true = c(0.04, 0.12, 0.25, 0.43, 0.63),
  p_resp = c(0.10, 0.20, 0.30, 0.45, 0.58),
  n_cohort = 10, cohort_size = 3, window = 1, accrual_rate = 3,
  n_trials = 200, seed = 1
)
print(oc)
print(oc, percent = TRUE)
```

---

print.backfill_oc_multi

Print Operating Characteristics of a Backfill Design Across Scenarios

---

Description

Display the result of sim_bf_boin_multi or sim_be_boin_multi as one table with seven rows per scenario.

Usage

```
## S3 method for class 'backfill_oc_multi'
print(
  x,
  digits = 1,
  percent = FALSE,
  kable = FALSE,
  kable_format = "pipe",
  ...
)
```

Arguments

x	An object of class backfill_oc_multi.
digits	Integer scalar. Number of decimal places. Defaults to 1.
percent	Logical scalar. Show patients, backfilled patients and DLTs as percentages of their totals. Defaults to FALSE.
kable	Logical scalar. Print only the table, formatted by knitr::kable(). Defaults to FALSE.
kable_format	Character scalar passed to knitr::kable() as format. Defaults to "pipe".
...	Further arguments, currently ignored.

Value

The object x, invisibly.

See Also

sim_bf_boin_multi, sim_be_boin_multi

Examples

```
scenarios <- list(
  list(name = "Rising", p_true = c(0.04, 0.12, 0.25, 0.43, 0.63),
       p_resp = c(0.10, 0.20, 0.30, 0.45, 0.58)),
  list(name = "Plateau", p_true = c(0.04, 0.12, 0.25, 0.43, 0.63),
       p_resp = c(0.30, 0.32, 0.35, 0.36, 0.36))
)
oc <- sim_bf_boin_multi(
  target = 0.25, scenarios = scenarios, n_cohort = 10, cohort_size = 3,
  window = 1, accrual_rate = 3, n_trials = 100, seed = 1
)
print(oc, percent = TRUE)
```

---

print.backfill_trials *Print Simulated Trials of a BOIN Design with Backfilling*

Description

Summarize the trials returned by bf_boin_simulate or be_boin_simulate: the design, the average numbers of patients, backfilled patients, DLTs and responses at each dose, the trial duration, the waiting of the dose escalation and the reasons for stopping.

Usage

```
## S3 method for class 'backfill_trials'
print(x, ...)
```

Arguments

x	An object of class backfill_trials.
...	Further arguments, currently ignored.

Value

The object x, invisibly.

See Also

bf_boin_simulate, be_boin_simulate

Examples

```
trials <- bf_boin_simulate(
  target = 0.25,
  p_true = c(0.04, 0.12, 0.25, 0.43, 0.63),
  p_resp = c(0.10, 0.20, 0.30, 0.45, 0.58),
  n_cohort = 10, cohort_size = 3, window = 1, accrual_rate = 3,
  n_trials = 100, seed = 1
)
print(trials)
```

---

print.boin_boundary *Print BOIN Decision Boundaries*

Description

Display the interval boundaries and the integer decision boundaries produced by boin_boundary.

Usage

```
## S3 method for class 'boin_boundary'
print(x, cohort_size = NULL, ...)
```

Arguments

x	An object of class boin_boundary.
cohort_size	Integer scalar or NULL. When supplied, only the sample sizes that are multiples of cohort_size are shown, which is how the boundaries are consulted during a trial with equally sized cohorts.
...	Further arguments, currently ignored.

Value

The object x, invisibly.

Examples

```
bd <- boin_boundary(target = 0.30, max_n = 18, extrasafe = TRUE)
print(bd, cohort_size = 3)
```

`print.boin_decision_table` *Print a BOIN Decision Table*

Description

Display the decision table produced by `boin_decision_table`, optionally restricted to the sample sizes reached at the end of a cohort.

Usage

```
## S3 method for class 'boin_decision_table'
print(x, cohort_size = NULL, ...)
```

Arguments

`x`	An object of class `boin_decision_table`.
`cohort_size`	Integer scalar or `NULL`. When supplied, only the sample sizes that are multiples of `cohort_size` are shown, which is how the table is consulted during a trial with equally sized cohorts.
`...`	Further arguments, currently ignored.

Value

The object `x`, invisibly.

See Also

`boin_decision_table`, `plot.boin_decision_table`

Examples

```
decisions <- boin_decision_table(target = 0.30, max_n = 18)

print(decisions, cohort_size = 3)
```

print.boin_oc *Print Operating Characteristics of a BOIN Design*

Description

Display the summary table produced by sim_boin.

Usage

```
## S3 method for class 'boin_oc'
print(
  x,
  digits = 1,
  percent = FALSE,
  kable = FALSE,
  kable_format = "pipe",
  ...
)
```

Arguments

x	An object of class boin_oc.
digits	Integer scalar. Number of decimal places. Defaults to 1.
percent	Logical scalar. Express the average number of patients and DLTs at each dose as a percentage of the trial total instead of a count. Defaults to FALSE.
kable	Logical scalar. Return the table through knitr::kable() rather than printing it directly, which is convenient inside a report. Requires the **knitr** package. Defaults to FALSE.
kable_format	Character scalar passed to knitr::kable(), for example "pipe", "html" or "latex". Defaults to "pipe".
...	Further arguments, currently ignored.

Value

The object x, invisibly.

Examples

```
oc <- sim_boin(
  target = 0.30,
  p_true = c(0.05, 0.15, 0.30, 0.45, 0.60),
  n_cohort = 10,
  cohort_size = 3,
  n_trials = 200,
  seed = 123
)
```

```
print(oc)
print(oc, percent = TRUE)
```

`print.boin_oc_multi` *Print Operating Characteristics Across Scenarios*

Description

Display the combined summary table produced by `sim_boin_multi`.

Usage

```
## S3 method for class 'boin_oc_multi'
print(
  x,
  digits = 1,
  percent = FALSE,
  kable = FALSE,
  kable_format = "pipe",
  ...
)
```

Arguments

`x`	An object of class `boin_oc_multi`.
`digits`	Integer scalar. Number of decimal places. Defaults to 1.
`percent`	Logical scalar. Express the average number of patients and DLTs at each dose as a percentage of the trial total instead of a count. Defaults to `FALSE`.
`kable`	Logical scalar. Return the table through `knitr::kable()` rather than printing it directly. Requires the **knitr** package. Defaults to `FALSE`.
`kable_format`	Character scalar passed to `knitr::kable()`, for example `"pipe"`, `"html"` or `"latex"`. Defaults to `"pipe"`.
`...`	Further arguments, currently ignored.

Value

The object `x`, invisibly.

Examples

```
oc <- sim_boin_multi(
  target = 0.30,
  scenarios = list(
    list(name = "MTD at dose 3", p_true = c(0.05, 0.15, 0.30, 0.45, 0.60)),
    list(name = "All toxic",     p_true = c(0.35, 0.45, 0.55, 0.65, 0.75))
  ),
  n_cohort = 10,
  cohort_size = 3,
  n_trials = 200,
  seed = 123
)

print(oc)
print(oc, percent = TRUE)
```

---

print.boin_trials *Print Simulated BOIN Trials*

Description

Show a compact summary of the object returned by `boin_simulate` instead of printing the full trial matrices.

Usage

```
## S3 method for class 'boin_trials'
print(x, ...)
```

Arguments

`x`	An object of class `boin_trials`.
`...`	Further arguments, currently ignored.

Value

The object `x`, invisibly.

Examples

```
trials <- boin_simulate(
  target = 0.30,
  p_true = c(0.05, 0.15, 0.30, 0.45, 0.60),
  n_cohort = 10,
  cohort_size = 3,
  n_trials = 100,
  seed = 1
)
```

```
print(trials)
```

`print.comb_boin_next` *Print the Next Dose Combination of a BOIN Drug-Combination Trial*

Description

Display the decision returned by `comb_boin_next`: the data at the current combination, the boundaries that apply to it, the decision, the candidates that were compared and the combinations eliminated so far.

Usage

```
## S3 method for class 'comb_boin_next'
print(x, ...)
```

Arguments

`x`	An object of class `comb_boin_next`.
`...`	Further arguments, currently ignored.

Value

The object `x`, invisibly.

Examples

```
n_pts <- matrix(c(3, 0, 0,
                  0, 0, 0), nrow = 2, byrow = TRUE)
n_tox <- matrix(0, nrow = 2, ncol = 3)

print(comb_boin_next(target = 0.30, n_pts = n_pts, n_tox = n_tox,
                     current = c(1, 1)))
```

`print.comb_boin_oc` *Print Operating Characteristics of the BOIN Drug-Combination Design*

Description

Display the summary produced by `sim_comb_boin`: the true DLT rates, the selection percentages and the average numbers of patients and DLTs at every dose combination, followed by the selection of the true MTD and the exposure to combinations above the overdose cutoff.

Usage

```
## S3 method for class 'comb_boin_oc'
print(x, digits = 1, kable = FALSE, kable_format = "pipe", ...)
```

Arguments

x	An object of class comb_boin_oc.
digits	Integer scalar. Number of decimal places. Defaults to 1.
kable	Logical scalar. Return the table through knitr::kable() rather than printing it directly, which is convenient inside a report. Requires the **knitr** package. Defaults to FALSE.
kable_format	Character scalar passed to knitr::kable(), for example "pipe", "html" or "latex". Defaults to "pipe".
...	Further arguments, currently ignored.

Value

The object x, invisibly.

Examples

```
p_true <- matrix(c(0.02, 0.04, 0.08, 0.14,
                   0.08, 0.25, 0.42, 0.48,
                   0.25, 0.45, 0.50, 0.60), nrow = 3, byrow = TRUE)

oc <- sim_comb_boin(
  target = 0.25,
  p_true = p_true,
  n_cohort = 16,
  cohort_size = 3,
  n_trials = 200,
  seed = 123
)

print(oc)
print(oc, digits = 2)
```

---

print.comb_boin_oc_multi

Print Operating Characteristics of the BOIN Drug-Combination Design Across Scenarios

---

Description

Display the table produced by sim_comb_boin_multi, one row per scenario. The full results of a scenario are printed by printing its element of x$results.

Usage

```
## S3 method for class 'comb_boin_oc_multi'
print(x, digits = 1, kable = FALSE, kable_format = "pipe", ...)
```

Arguments

x	An object of class comb_boin_oc_multi.
digits	Integer scalar. Number of decimal places. Defaults to 1.
kable	Logical scalar. Return the table through knitr::kable() rather than printing it directly. Requires the **knitr** package. Defaults to FALSE.
kable_format	Character scalar passed to knitr::kable(), for example "pipe", "html" or "latex". Defaults to "pipe".
...	Further arguments, currently ignored.

Value

The object x, invisibly.

Examples

```
oc <- sim_comb_boin_multi(
  target = 0.30,
  scenarios = list(
    "MTD in the middle" = matrix(c(0.10, 0.20, 0.30,
                                   0.20, 0.30, 0.45), 2, byrow = TRUE),
    "All toxic"         = matrix(c(0.40, 0.50, 0.60,
                                   0.50, 0.60, 0.70), 2, byrow = TRUE)
  ),
  n_cohort = 10,
  cohort_size = 3,
  n_trials = 200,
  seed = 123
)

print(oc)
```

---

print.comb_boin_trials

Print Simulated BOIN Drug-Combination Trials

Description

Show a compact summary of the object returned by comb_boin_simulate instead of printing the full trial arrays.

Usage

```
## S3 method for class 'comb_boin_trials'
print(x, ...)
```

Arguments

x	An object of class comb_boin_trials.
...	Further arguments, currently ignored.

Value

The object x, invisibly.

Examples

```
p_true <- matrix(c(0.05, 0.15, 0.30,
                   0.15, 0.30, 0.45), nrow = 2, byrow = TRUE)

trials <- comb_boin_simulate(
  target = 0.30,
  p_true = p_true,
  n_cohort = 10,
  cohort_size = 3,
  n_trials = 100,
  seed = 1
)
print(trials)
```

---

print.decision_table_3p3

Print a 3+3 Decision Table

Description

Display the decision table produced by decision_table_3p3: one row per number of patients and range of DLTs sharing a decision, for dose escalation and, with mtd_rule = "expand", for the search for the MTD, followed by what happens once escalation stops.

Usage

```
## S3 method for class 'decision_table_3p3'
print(x, ...)
```

Arguments

x	An object of class decision_table_3p3.
...	Further arguments, currently ignored.

Details

A data frame that no longer has the columns or attributes of a decision table, for example after selecting columns, is printed as a plain data frame.

Value

The object `x`, invisibly.

See Also

`decision_table_3p3`, `plot.decision_table_3p3`

Examples

```
print(decision_table_3p3())

print(decision_table_3p3(mtd_rule = "expand"))
```

`print.oc_3p3` *Print Operating Characteristics of the 3+3 Design*

Description

Display the summary table produced by `oc_3p3` or `sim_3p3`. The layout matches `print.boin_oc`, so a 3+3 table can be read beside a BOIN one.

Usage

```
## S3 method for class 'oc_3p3'
print(
  x,
  digits = 1,
  percent = FALSE,
  kable = FALSE,
  kable_format = "pipe",
  ...
)
```

Arguments

`x`	An object of class `oc_3p3`.
`digits`	Integer scalar. Number of decimal places. Defaults to 1.
`percent`	Logical scalar. Express the average number of patients and DLTs at each dose as a percentage of the trial total instead of a count. Defaults to `FALSE`.

kable	Logical scalar. Return the table through knitr::kable() rather than printing it directly, which is convenient inside a report. Requires the **knitr** package. Defaults to FALSE.
kable_format	Character scalar passed to knitr::kable(), for example "pipe", "html" or "latex". Defaults to "pipe".
...	Further arguments, currently ignored.

Value

The object x, invisibly.

See Also

oc_3p3, sim_3p3

Examples

```
oc <- oc_3p3(p_true = c(0.30, 0.48, 0.67))

print(oc)
print(oc, percent = TRUE)
```

---

print.tite_boin_decision_table

Print a TITE-BOIN Decision Table

Description

Display the decision table produced by tite_boin_decision_table in the layout of the published tables: one row per number of patients treated, number of DLTs and range of pending patients sharing a decision, with the boundaries on the follow-up statistic in the Escalate, Stay and De-escalate columns.

Usage

```
## S3 method for class 'tite_boin_decision_table'
print(x, cohort_size = NULL, digits = 2, ...)
```

Arguments

x	An object of class tite_boin_decision_table.
cohort_size	Integer scalar or NULL. When supplied, only the sample sizes that are multiples of cohort_size are shown, which is how the published tables are laid out.
digits	Integer scalar. Number of decimal places of the boundaries. Defaults to 2, as in the published tables.
...	Further arguments, currently ignored.

Details

A boundary shown as ">= 2.15" under Escalate means that the dose is escalated when the follow-up statistic is at least 2.15; the statistic is STFT for `method = "imputation"` and ESS for `method = "ess"`. "Suspend if >= 4.23" under Escalate marks an escalation blocked by `min_completed`. When the table was built with a positive `min_follow_up`, an escalation with pending patients carries the condition on MF, for example ">= 0.22 & MF >= 0.25"; accrual is suspended when the condition on the follow-up statistic holds but the one on MF does not.

A data frame that no longer has the columns or attributes of a decision table, for example after selecting columns, is printed as a plain data frame.

Value

The object `x`, invisibly.

See Also

`tite_boin_decision_table`, `plot.tite_boin_decision_table`

Examples

```
decisions <- tite_boin_decision_table(target = 0.3, max_n = 15)

# Table S1 of Yuan et al. (2018)
print(decisions, cohort_size = 3)

# Boundaries with three decimal places
print(decisions[decisions$n == 9, ], digits = 3)
```

---

print.tite_boin_oc *Print Operating Characteristics of a TITE-BOIN Design*

Description

Display the summary produced by `sim_tite_boin`: the design, the table of `print.boin_oc`, the exposure to doses above the overdose cutoff, and the duration of the trials and the suspensions of accrual.

Usage

```
## S3 method for class 'tite_boin_oc'
print(
  x,
  digits = 1,
  percent = FALSE,
  kable = FALSE,
  kable_format = "pipe",
  ...
)
```

Arguments

x	An object of class tite_boin_oc.
digits	Integer scalar. Number of decimal places. Defaults to 1.
percent	Logical scalar. Express the average number of patients and DLTs at each dose as a percentage of the trial total instead of a count. Defaults to FALSE.
kable	Logical scalar. Return the table through knitr::kable() rather than printing it directly, which is convenient inside a report. Requires the **knitr** package. Defaults to FALSE.
kable_format	Character scalar passed to knitr::kable(), for example "pipe", "html" or "latex". Defaults to "pipe".
...	Further arguments, currently ignored.

Value

The object x, invisibly.

Examples

```
oc <- sim_tite_boin(
  target = 0.30,
  p_true = c(0.05, 0.15, 0.30, 0.45, 0.60),
  n_cohort = 10,
  cohort_size = 3,
  window = 3,
  accrual_rate = 2,
  n_trials = 200,
  seed = 123
)

print(oc)
print(oc, percent = TRUE)
```

---

print.tite_boin_oc_multi *Print Operating Characteristics of a TITE-BOIN Design Across Scenarios*

Description

Display the table produced by sim_tite_boin_multi, one block of rows per scenario.

Usage

```
## S3 method for class 'tite_boin_oc_multi'
print(
  x,
  digits = 1,
  percent = FALSE,
  kable = FALSE,
  kable_format = "pipe",
  ...
)
```

Arguments

x	An object of class tite_boin_oc_multi.
digits	Integer scalar. Number of decimal places. Defaults to 1.
percent	Logical scalar. Express the average number of patients and DLTs at each dose as a percentage of the trial total instead of a count. Defaults to FALSE.
kable	Logical scalar. Return the table through knitr::kable(). Requires the **knitr** package. Defaults to FALSE.
kable_format	Character scalar passed to knitr::kable(). Defaults to "pipe".
...	Further arguments, currently ignored.

Value

The object x, invisibly.

Examples

```
scenarios <- list(
  list(name = "MTD at dose 3", p_true = c(0.05, 0.15, 0.30, 0.45, 0.60)),
  list(name = "All toxic",     p_true = c(0.35, 0.45, 0.55, 0.65, 0.75))
)

oc <- sim_tite_boin_multi(
  target = 0.30,
  scenarios = scenarios,
  n_cohort = 10,
  cohort_size = 3,
  window = 3,
  accrual_rate = 2,
  n_trials = 200,
  seed = 123
)

print(oc)
print(oc, percent = TRUE)
```

print.tite_boin_trials *Print Simulated TITE-BOIN Trials*

Description

Show a compact summary of the object returned by `tite_boin_simulate` instead of printing the full trial matrices.

Usage

```
## S3 method for class 'tite_boin_trials'
print(x, ...)
```

Arguments

`x`	An object of class `tite_boin_trials`.
`...`	Further arguments, currently ignored.

Value

The object `x`, invisibly.

Examples

```
trials <- tite_boin_simulate(
  target = 0.30,
  p_true = c(0.05, 0.15, 0.30, 0.45, 0.60),
  n_cohort = 10,
  cohort_size = 3,
  window = 3,
  accrual_rate = 2,
  n_trials = 100,
  seed = 1
)
print(trials)
```

sim_3p3 *Simulate the 3+3 Design*

Description

Simulate the traditional 3+3 dose-finding design and summarize the same quantities that `oc_3p3` computes exactly. The two are meant to be used together: `oc_3p3()` gives the answer without Monte Carlo error, and this function confirms it from independent trials.

Usage

```
sim_3p3(
  p_true,
  n_trials = 10000,
  mtd_rule = c("previous", "expand"),
  start_dose = 1,
  overdose_cutoff = NULL,
  seed = 123
)
```

Arguments

p_true	Numeric vector. True DLT probability at each dose level, in increasing dose order.
n_trials	Integer scalar. Number of trials to simulate. Defaults to 10000.
mtd_rule	Character scalar, either "previous" or "expand". See oc_3p3. Defaults to "previous".
start_dose	Integer scalar. Dose level for the first cohort. Defaults to 1.
overdose_cutoff	Numeric scalar or NULL. Doses whose true DLT probability exceeds this value count as overdoses in the overdose component of the result. Defaults to NULL, which uses one third.
seed	Integer scalar or NULL. Random seed. The state of the calling session is restored on exit. Defaults to 123.

Details

The design is described under oc_3p3. Since the operating characteristics of the 3+3 design can be written in closed form, this function is a check on that closed form rather than the way to obtain the numbers. Simulation is the only option when a quantity is wanted that the closed form does not provide, such as the distribution of the sample size rather than its mean.

Value

An object of class oc_3p3, with the components described in oc_3p3. The method component is the number of trials rather than "exact".

See Also

oc_3p3, sim_boin

Examples

```
exact <- oc_3p3(p_true = c(0.30, 0.48, 0.67))
simulated <- sim_3p3(p_true = c(0.30, 0.48, 0.67), n_trials = 2000, seed = 1)

round(exact$sel_percent, 2)
round(simulated$sel_percent, 2)
```

`sim_be_boin` *Operating Characteristics of a BE-BOIN Design*

Description

Simulate the BOIN design with backfilling and late-onset toxicity (BE-BOIN) of Chen et al. (2026) many times under one scenario of DLT and response probabilities, and summarize how often each dose is selected as the MTD, how many patients are treated and backfilled at each dose, how many DLTs and responses occur, and how long the trials last.

Usage

```
sim_be_boin(
  target,
  p_true,
  p_resp,
  n_cohort,
  cohort_size,
  window,
  accrual_rate,
  n_cap = 12,
  backfill_dose = c("highest", "lowest"),
  conflict_dose = c("lowest", "highest"),
  no_slot = c("wait", "leave"),
  accrual = c("exponential", "uniform", "fixed"),
  dlt_time = c("weibull", "uniform"),
  late_fraction = 0.5,
  resp_window = window,
  resp_late_fraction = 0.5,
  resp_cor = 0,
  max_pending_ratio = 0.49,
  min_completed = 0,
  min_follow_up = 0.25,
  n_trials = 10000,
  start_dose = 1,
  n_earlystop = 18,
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  bound_mtd = FALSE,
  mtd_max_estimate = NULL,
  min_mtd_sample = 1,
  overdose_cutoff = NULL,
  n_earlystop_rule = c("with_stay", "simple"),
  keep_trials = FALSE,
```

```
  verbose = FALSE,
  seed = 123
)
```

Arguments

target	Numeric scalar. Target DLT probability, for example 0.25.
p_true	Numeric vector. True DLT probability at each dose level, in increasing dose order. The probability of a DLT within the assessment window.
p_resp	Numeric vector of the same length as p_true. True probability of a response within the response window at each dose. A response at or below a dose is what makes the dose eligible for backfilling.
n_cohort	Integer scalar. Number of cohorts in the dose escalation.
cohort_size	Integer scalar or vector. Number of patients per cohort of the dose escalation. A scalar is used for every cohort. A vector shorter than n_cohort is padded with its last element and a longer one is truncated.
window	Numeric scalar. Length of the DLT assessment window.
accrual_rate	Numeric scalar. Average number of patients arriving per unit of time, in the same unit as window.
n_cap	Integer scalar. A dose is no longer backfilled once this many patients, from the dose escalation and from backfilling together, have been treated at it. The dose escalation can still treat patients at the dose. Defaults to 12.
backfill_dose	Character scalar. Which dose receives a backfill patient when several are open: "highest" (the default, as in both articles) or "lowest".
conflict_dose	Character scalar. When the data of several backfilled doses conflict with those of the current dose, which of them anchors the pooled estimate: "lowest" (the default here, following Chen et al., 2026) or "highest" (the default of bf_boin_simulate, following Zhao et al., 2024).
no_slot	Character scalar. What happens to a patient who arrives when neither the dose escalation nor any backfill dose can take a patient: "wait" (the default), the patient waits and is treated as soon as one can, as in tite_boin_simulate; or "leave", the patient is turned away and the trial waits for the next arrival.
accrual	Character scalar. Distribution of the times between arrivals: "exponential" (the default) for a Poisson process, "uniform" for a uniform distribution on (0, 2 / accrual_rate), or "fixed" for arrivals exactly 1 / accrual_rate apart.
dlt_time	Character scalar. Distribution of the time to DLT within the window, and of the time to response within the response window: "weibull" (the default) or "uniform".
late_fraction	Numeric scalar strictly between 0 and 1. Share of the DLTs that occur in the second half of the assessment window under the Weibull distribution. Defaults to 0.5.
resp_window	Numeric scalar. Length of the window within which a response is observed. Defaults to window.

`resp_late_fraction`
: Numeric scalar strictly between 0 and 1. Share of the responses that occur in the second half of the response window under the Weibull distribution. Defaults to 0.5.

`resp_cor`
: Numeric scalar strictly between -1 and 1. Correlation of the Gaussian copula that links the DLT and the response of a patient; a positive value makes a response more likely in a patient with a DLT. Defaults to 0, independence.

`max_pending_ratio`
: Numeric scalar greater than 0 and at most 1. Rule 1 of the design: the escalation is suspended when the proportion of pending patients at the current dose exceeds this value, unless the dose is de-escalated whatever their outcomes. Defaults to 0.49, which suspends when fewer than 51 percent have completed the assessment, as in Chen et al. (2026). See `tite_boin_decision_table`.

`min_completed`
: Integer scalar. While some patients are pending, escalation requires at least this many patients with a completed assessment. Defaults to 0, which the design does not use. See `tite_boin_decision_table`.

`min_follow_up`
: Numeric scalar between 0 and 1. Rule 2 of the design: escalation requires every pending patient at the current dose to have been followed for at least this fraction of the window. Defaults to 0.25, as in Chen et al. (2026). See `tite_boin_decision_table`.

`n_trials`
: Integer scalar. Number of trials to simulate. Defaults to 10000.

`start_dose`
: Integer scalar. Dose level for the first cohort. Defaults to 1.

`n_earlystop`
: Integer scalar. The trial stops once this many patients, backfilled ones included, have been treated at the current dose and the design would stay there. Defaults to 18.

`p_saf`
: Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to `0.6 * target`.

`p_tox`
: Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to `1.4 * target`.

`cutoff_eli`
: Numeric scalar. Posterior probability cutoff for dose elimination. Defaults to 0.95.

`extrasafe`
: Logical scalar. Apply the stricter safety stopping rule at the lowest dose. Defaults to `FALSE`.

`offset`
: Numeric scalar between 0 and 0.5. Amount by which `cutoff_eli` is relaxed for the safety stopping rule. Defaults to 0.05.

`bound_mtd`
: Logical scalar. Require the isotonic estimate at the selected dose to be at or below the de-escalation boundary. Defaults to `FALSE`.

`mtd_max_estimate`
: Numeric scalar or `NULL`. Largest isotonic estimate a dose may have and still be selected as the MTD. See `sim_boin`.

`min_mtd_sample`
: Integer scalar. Smallest number of patients a dose must have received to be eligible as the MTD. Defaults to 1.

`overdose_cutoff`
: Numeric scalar or `NULL`. Doses whose true DLT probability exceeds this value count as overdoses in the `overdose` component of the result. Defaults to `NULL`, which uses `target`.

`n_earlystop_rule`
: Character scalar, either `"with_stay"` or `"simple"`. See `boin_simulate`.

`keep_trials`
: Logical scalar. Keep the full trial by trial data in the result. Defaults to `FALSE`.

`verbose`
: Logical scalar. Report progress while the simulation runs. Defaults to `FALSE`.

`seed`
: Integer scalar or `NULL`. Random seed. When `NULL` the current state of the random number generator is used. The state of the calling session is restored on exit. Defaults to 123.

Details

The trials are simulated by `be_boin_simulate`, whose details describe the conduct of a trial. The MTD is selected as in `sim_bf_boin`.

When no dose is ever opened for backfilling, for example with every `p_resp` equal to 0, and with `no_slot = "wait"`, every result that the two functions share agrees with that of `sim_tite_boin` with `method = "imputation"`, the same suspension rules and the same seed.

Only the MTD is selected. The second stage of Chen et al. (2026), which randomizes patients between doses to choose the optimal biological dose, is not part of the simulation.

Value

An object of class `c("be_boin_oc", "backfill_oc", "boin_oc")`, with the components described in `sim_bf_boin`.

References

Chen, K., Zhao, Y., Takeda, K. and Yuan, Y. (2026). BE-BOIN: A Dose Optimization Design Accommodating Backfill and Late-Onset Toxicity. Therapeutic Innovation and Regulatory Science. doi:10.1007/s43441026009940

See Also

`be_boin_simulate`, `sim_be_boin_multi`, `sim_bf_boin`, `sim_tite_boin`

Examples

```
# A three month window with two patients a month
oc <- sim_be_boin(
  target = 0.25,
  p_true = c(0.04, 0.12, 0.25, 0.43, 0.63),
  p_resp = c(0.10, 0.20, 0.30, 0.45, 0.58),
  n_cohort = 10,
  cohort_size = 3,
  window = 3,
  accrual_rate = 2,
  n_earlystop = 9,
  n_trials = 500,
  seed = 123
)
oc
```

`sim_be_boin_multi` *Operating Characteristics of a BE-BOIN Design Across Several Scenarios*

Description

Run `sim_be_boin` under a list of scenarios of DLT and response probabilities that share the same design, and collect the results into one table.

Usage

```
sim_be_boin_multi(
  target,
  scenarios,
  n_cohort,
  cohort_size,
  window,
  accrual_rate,
  n_cap = 12,
  backfill_dose = c("highest", "lowest"),
  conflict_dose = c("lowest", "highest"),
  no_slot = c("wait", "leave"),
  accrual = c("exponential", "uniform", "fixed"),
  dlt_time = c("weibull", "uniform"),
  late_fraction = 0.5,
  resp_window = window,
  resp_late_fraction = 0.5,
  resp_cor = 0,
  max_pending_ratio = 0.49,
  min_completed = 0,
  min_follow_up = 0.25,
  n_trials = 10000,
  start_dose = 1,
  n_earlystop = 18,
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  bound_mtd = FALSE,
  mtd_max_estimate = NULL,
  min_mtd_sample = 1,
  overdose_cutoff = NULL,
  n_earlystop_rule = c("with_stay", "simple"),
  keep_trials = FALSE,
  verbose = FALSE,
  seed = 123
```

)

Arguments

target	Numeric scalar. Target DLT probability, for example 0.25.
scenarios	A list of scenarios. Each element is a list with a p_true vector of true DLT probabilities, a p_resp vector of true response probabilities of the same length and optionally a name; without a name the list names are used, or "Scenario 1", "Scenario 2" and so on. Every scenario must have the same number of doses.
n_cohort	Integer scalar. Number of cohorts in the dose escalation.
cohort_size	Integer scalar or vector. Number of patients per cohort of the dose escalation. A scalar is used for every cohort. A vector shorter than n_cohort is padded with its last element and a longer one is truncated.
window	Numeric scalar. Length of the DLT assessment window.
accrual_rate	Numeric scalar. Average number of patients arriving per unit of time, in the same unit as window.
n_cap	Integer scalar. A dose is no longer backfilled once this many patients, from the dose escalation and from backfilling together, have been treated at it. The dose escalation can still treat patients at the dose. Defaults to 12.
backfill_dose	Character scalar. Which dose receives a backfill patient when several are open: "highest" (the default, as in both articles) or "lowest".
conflict_dose	Character scalar. When the data of several backfilled doses conflict with those of the current dose, which of them anchors the pooled estimate: "lowest" (the default here, following Chen et al., 2026) or "highest" (the default of bf_boin_simulate, following Zhao et al., 2024).
no_slot	Character scalar. What happens to a patient who arrives when neither the dose escalation nor any backfill dose can take a patient: "wait" (the default), the patient waits and is treated as soon as one can, as in tite_boin_simulate; or "leave", the patient is turned away and the trial waits for the next arrival.
accrual	Character scalar. Distribution of the times between arrivals: "exponential" (the default) for a Poisson process, "uniform" for a uniform distribution on (0, 2 / accrual_rate), or "fixed" for arrivals exactly 1 / accrual_rate apart.
dlt_time	Character scalar. Distribution of the time to DLT within the window, and of the time to response within the response window: "weibull" (the default) or "uniform".
late_fraction	Numeric scalar strictly between 0 and 1. Share of the DLTs that occur in the second half of the assessment window under the Weibull distribution. Defaults to 0.5.
resp_window	Numeric scalar. Length of the window within which a response is observed. Defaults to window.
resp_late_fraction	Numeric scalar strictly between 0 and 1. Share of the responses that occur in the second half of the response window under the Weibull distribution. Defaults to 0.5.

`resp_cor` Numeric scalar strictly between -1 and 1. Correlation of the Gaussian copula that links the DLT and the response of a patient; a positive value makes a response more likely in a patient with a DLT. Defaults to 0, independence.

`max_pending_ratio` Numeric scalar greater than 0 and at most 1. Rule 1 of the design: the escalation is suspended when the proportion of pending patients at the current dose exceeds this value, unless the dose is de-escalated whatever their outcomes. Defaults to 0.49, which suspends when fewer than 51 percent have completed the assessment, as in Chen et al. (2026). See `tite_boin_decision_table`.

`min_completed` Integer scalar. While some patients are pending, escalation requires at least this many patients with a completed assessment. Defaults to 0, which the design does not use. See `tite_boin_decision_table`.

`min_follow_up` Numeric scalar between 0 and 1. Rule 2 of the design: escalation requires every pending patient at the current dose to have been followed for at least this fraction of the window. Defaults to 0.25, as in Chen et al. (2026). See `tite_boin_decision_table`.

`n_trials` Integer scalar. Number of trials to simulate. Defaults to 10000.

`start_dose` Integer scalar. Dose level for the first cohort. Defaults to 1.

`n_earlystop` Integer scalar. The trial stops once this many patients, backfilled ones included, have been treated at the current dose and the design would stay there. Defaults to 18.

`p_saf` Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to `0.6 * target`.

`p_tox` Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to `1.4 * target`.

`cutoff_eli` Numeric scalar. Posterior probability cutoff for dose elimination. Defaults to 0.95.

`extrasafe` Logical scalar. Apply the stricter safety stopping rule at the lowest dose. Defaults to `FALSE`.

`offset` Numeric scalar between 0 and 0.5. Amount by which `cutoff_eli` is relaxed for the safety stopping rule. Defaults to 0.05.

`bound_mtd` Logical scalar. Require the isotonic estimate at the selected dose to be at or below the de-escalation boundary. Defaults to `FALSE`.

`mtd_max_estimate` Numeric scalar or `NULL`. Largest isotonic estimate a dose may have and still be selected as the MTD. See `sim_boin`.

`min_mtd_sample` Integer scalar. Smallest number of patients a dose must have received to be eligible as the MTD. Defaults to 1.

`overdose_cutoff` Numeric scalar or `NULL`. Doses whose true DLT probability exceeds this value count as overdoses in the `overdose` component of the result. Defaults to `NULL`, which uses `target`.

`n_earlystop_rule` Character scalar, either `"with_stay"` or `"simple"`. See `boin_simulate`.

keep_trials	Logical scalar. Keep the full trial by trial data in the result. Defaults to FALSE.
verbose	Logical scalar. Report progress while the simulation runs. Defaults to FALSE.
seed	Integer scalar or NULL. Random seed, applied to every scenario, so that a scenario simulated here matches the same scenario simulated on its own with sim_be_boin. Defaults to 123.

Value

An object of class c("be_boin_oc_multi", "backfill_oc_multi", "boin_oc_multi"), with the components described in sim_bf_boin_multi; results holds be_boin_oc objects.

See Also

sim_be_boin, sim_bf_boin_multi

Examples

```
scenarios <- list(
  list(name = "MTD at dose 3", p_true = c(0.04, 0.12, 0.25, 0.43, 0.63),
       p_resp = c(0.10, 0.20, 0.30, 0.45, 0.58)),
  list(name = "Plateau", p_true = c(0.04, 0.12, 0.25, 0.43, 0.63),
       p_resp = c(0.30, 0.32, 0.35, 0.36, 0.36))
)

oc <- sim_be_boin_multi(
  target = 0.25,
  scenarios = scenarios,
  n_cohort = 10,
  cohort_size = 3,
  window = 3,
  accrual_rate = 2,
  n_earlystop = 9,
  n_trials = 200,
  seed = 123
)
oc
```

---

sim_bf_boin *Operating Characteristics of a BF-BOIN Design*

Description

Simulate the BOIN design with backfilling (BF-BOIN) of Zhao et al. (2024) many times under one scenario of DLT and response probabilities, and summarize how often each dose is selected as the MTD, how many patients are treated and backfilled at each dose, how many DLTs and responses occur, and how long the trials last.

Usage

```
sim_bf_boin(
  target,
  p_true,
  p_resp,
  n_cohort,
  cohort_size,
  window,
  accrual_rate,
  n_cap = 12,
  backfill_dose = c("highest", "lowest"),
  conflict_dose = c("highest", "lowest"),
  no_slot = c("wait", "leave"),
  accrual = c("exponential", "uniform", "fixed"),
  dlt_time = c("weibull", "uniform"),
  late_fraction = 0.5,
  resp_window = window,
  resp_late_fraction = 0.5,
  resp_cor = 0,
  n_trials = 10000,
  start_dose = 1,
  n_earlystop = 18,
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  stay_on_1_of_3 = FALSE,
  bound_mtd = FALSE,
  mtd_max_estimate = NULL,
  min_mtd_sample = 1,
  overdose_cutoff = NULL,
  n_earlystop_rule = c("with_stay", "simple"),
  keep_trials = FALSE,
  verbose = FALSE,
  seed = 123
)
```

Arguments

target	Numeric scalar. Target DLT probability, for example 0.25.
p_true	Numeric vector. True DLT probability at each dose level, in increasing dose order. The probability of a DLT within the assessment window.
p_resp	Numeric vector of the same length as p_true. True probability of a response within the response window at each dose. A response at or below a dose is what makes the dose eligible for backfilling.
n_cohort	Integer scalar. Number of cohorts in the dose escalation.

`cohort_size` Integer scalar or vector. Number of patients per cohort of the dose escalation. A scalar is used for every cohort. A vector shorter than `n_cohort` is padded with its last element and a longer one is truncated.

`window` Numeric scalar. Length of the DLT assessment window.

`accrual_rate` Numeric scalar. Average number of patients arriving per unit of time, in the same unit as `window`.

`n_cap` Integer scalar. A dose is no longer backfilled once this many patients, from the dose escalation and from backfilling together, have been treated at it. The dose escalation can still treat patients at the dose. Defaults to 12.

`backfill_dose` Character scalar. Which dose receives a backfill patient when several are open: `"highest"` (the default, as in both articles) or `"lowest"`.

`conflict_dose` Character scalar. When the data of several backfilled doses conflict with those of the current dose, which of them anchors the pooled estimate: `"highest"` (the default here, following Zhao et al., 2024) or `"lowest"` (the default of `be_boin_simulate`, following Chen et al., 2026). See the details.

`no_slot` Character scalar. What happens to a patient who arrives when neither the dose escalation nor any backfill dose can take a patient: `"wait"` (the default), the patient waits and is treated as soon as one can, as in `tite_boin_simulate`; or `"leave"`, the patient is turned away and the trial waits for the next arrival.

`accrual` Character scalar. Distribution of the times between arrivals: `"exponential"` (the default) for a Poisson process, `"uniform"` for a uniform distribution on `(0, 2 / accrual_rate)`, or `"fixed"` for arrivals exactly `1 / accrual_rate` apart.

`dlt_time` Character scalar. Distribution of the time to DLT within the window, and of the time to response within the response window: `"weibull"` (the default) or `"uniform"`.

`late_fraction` Numeric scalar strictly between 0 and 1. Share of the DLTs that occur in the second half of the assessment window under the Weibull distribution. Defaults to 0.5.

`resp_window` Numeric scalar. Length of the window within which a response is observed. Defaults to `window`.

`resp_late_fraction` Numeric scalar strictly between 0 and 1. Share of the responses that occur in the second half of the response window under the Weibull distribution. Defaults to 0.5.

`resp_cor` Numeric scalar strictly between -1 and 1. Correlation of the Gaussian copula that links the DLT and the response of a patient; a positive value makes a response more likely in a patient with a DLT. Defaults to 0, independence.

`n_trials` Integer scalar. Number of trials to simulate. Defaults to 10000.

`start_dose` Integer scalar. Dose level for the first cohort. Defaults to 1.

`n_earlystop` Integer scalar. The trial stops once this many patients, backfilled ones included, have been treated at the current dose and the design would stay there. Defaults to 18.

`p_saf` Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to `0.6 * target`.

p_tox — Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to 1.4 * target.

cutoff_eli — Numeric scalar. Posterior probability cutoff for dose elimination. Defaults to 0.95.

extrasafe — Logical scalar. Apply the stricter safety stopping rule at the lowest dose. Defaults to FALSE.

offset — Numeric scalar between 0 and 0.5. Amount by which cutoff_eli is relaxed for the safety stopping rule. Defaults to 0.05.

stay_on_1_of_3 — Logical scalar. When TRUE, one DLT out of three patients leads to staying rather than de-escalating, in every use of the decision table. Zhao et al. (2024) use this modification in their main simulation. Defaults to FALSE. See boin_boundary.

bound_mtd — Logical scalar. Require the isotonic estimate at the selected dose to be at or below the de-escalation boundary. Defaults to FALSE.

mtd_max_estimate — Numeric scalar or NULL. Largest isotonic estimate a dose may have and still be selected as the MTD. See sim_boin.

min_mtd_sample — Integer scalar. Smallest number of patients a dose must have received to be eligible as the MTD. Defaults to 1.

overdose_cutoff — Numeric scalar or NULL. Doses whose true DLT probability exceeds this value count as overdoses in the overdose component of the result. Defaults to NULL, which uses target.

n_earlystop_rule — Character scalar, either "with_stay" or "simple". See boin_simulate.

keep_trials — Logical scalar. Keep the full trial by trial data in the result. Defaults to FALSE.

verbose — Logical scalar. Report progress while the simulation runs. Defaults to FALSE.

seed — Integer scalar or NULL. Random seed. When NULL the current state of the random number generator is used. The state of the calling session is restored on exit. Defaults to 123.

Details

The trials are simulated by bf_boin_simulate, whose details describe the conduct of a trial. Once every patient has completed the assessment the MTD is selected from the data of all patients, backfilled ones included, exactly as in sim_boin, and a trial stopped for safety selects no dose. The overdose cutoffs of 60 and 80 percent refer to the size of the dose escalation, n_cohort cohorts.

When no dose is ever opened for backfilling, for example with every p_resp equal to 0, the selection and the patient counts agree with those of sim_boin under the same seed.

To reproduce Table 4 of Zhao et al. (2024), use target = 0.25, n_cohort = 10, cohort_size = 3, window = 1, accrual_rate = 3, n_earlystop = 9, stay_on_1_of_3 = TRUE, accrual = "uniform" and no_slot = "leave".

Value

An object of class `c("bf_boin_oc", "backfill_oc", "boin_oc")`, which has every component of the result of `sim_boin`, so that code written for `sim_boin()` also reads it, counting backfilled patients among the patients treated, and in addition

`n_bf_dose` Average number of backfilled patients at each dose.

`total_n_bf` Average number of backfilled patients per trial.

`pct_trials_backfill` Percentage of trials with at least one backfilled patient.

`n_resp_dose` Average number of patients with a response at each dose.

`total_n_resp` Average number of patients with a response per trial.

`duration_mean` Average trial duration, from the first arrival until every enrolled patient has completed the DLT assessment, in the time unit of `window` and `accrual_rate`.

`duration_sd` Standard deviation of the trial duration.

`pct_trials_suspended` Percentage of trials in which the dose escalation had to wait at least once.

`avg_n_suspensions` Average number of decisions per trial at which the dose escalation had to wait.

`avg_time_suspended` Average total time per trial during which the dose escalation waited.

`avg_n_turned_away` Average number of patients turned away per trial, which is zero unless `no_slot = "leave"`.

`p_resp` The true response probabilities.

The `trials` component, when kept, holds the trial by trial data of `bf_boin_simulate` together with the selected MTD.

References

Zhao, Y., Yuan, Y., Korn, E. L. and Freidlin, B. (2024). Backfilling Patients in Phase I Dose-Escalation Trials Using Bayesian Optimal Interval Design (BOIN). Clinical Cancer Research, 30(4), 673-679.

See Also

`bf_boin_simulate`, `sim_bf_boin_multi`, `sim_be_boin`, `sim_boin`

Examples

```
# Scenario 3 of Zhao et al. (2024)
oc <- sim_bf_boin(
  target = 0.25,
  p_true = c(0.04, 0.12, 0.25, 0.43, 0.63),
  p_resp = c(0.10, 0.20, 0.30, 0.45, 0.58),
  n_cohort = 10,
  cohort_size = 3,
```

```
  window = 1,
  accrual_rate = 3,
  n_earlystop = 9,
  stay_on_1_of_3 = TRUE,
  accrual = "uniform",
  no_slot = "leave",
  n_trials = 500,
  seed = 123
)
oc
```

`sim_bf_boin_multi` *Operating Characteristics of a BF-BOIN Design Across Several Scenarios*

Description

Run `sim_bf_boin` under a list of scenarios of DLT and response probabilities that share the same design, and collect the results into one table.

Usage

```
sim_bf_boin_multi(
  target,
  scenarios,
  n_cohort,
  cohort_size,
  window,
  accrual_rate,
  n_cap = 12,
  backfill_dose = c("highest", "lowest"),
  conflict_dose = c("highest", "lowest"),
  no_slot = c("wait", "leave"),
  accrual = c("exponential", "uniform", "fixed"),
  dlt_time = c("weibull", "uniform"),
  late_fraction = 0.5,
  resp_window = window,
  resp_late_fraction = 0.5,
  resp_cor = 0,
  n_trials = 10000,
  start_dose = 1,
  n_earlystop = 18,
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
```

```
  stay_on_1_of_3 = FALSE,
  bound_mtd = FALSE,
  mtd_max_estimate = NULL,
  min_mtd_sample = 1,
  overdose_cutoff = NULL,
  n_earlystop_rule = c("with_stay", "simple"),
  keep_trials = FALSE,
  verbose = FALSE,
  seed = 123
)
```

Arguments

target	Numeric scalar. Target DLT probability, for example 0.25.
scenarios	A list of scenarios. Each element is a list with a p_true vector of true DLT probabilities, a p_resp vector of true response probabilities of the same length and optionally a name; without a name the list names are used, or "Scenario 1", "Scenario 2" and so on. Every scenario must have the same number of doses.
n_cohort	Integer scalar. Number of cohorts in the dose escalation.
cohort_size	Integer scalar or vector. Number of patients per cohort of the dose escalation. A scalar is used for every cohort. A vector shorter than n_cohort is padded with its last element and a longer one is truncated.
window	Numeric scalar. Length of the DLT assessment window.
accrual_rate	Numeric scalar. Average number of patients arriving per unit of time, in the same unit as window.
n_cap	Integer scalar. A dose is no longer backfilled once this many patients, from the dose escalation and from backfilling together, have been treated at it. The dose escalation can still treat patients at the dose. Defaults to 12.
backfill_dose	Character scalar. Which dose receives a backfill patient when several are open: "highest" (the default, as in both articles) or "lowest".
conflict_dose	Character scalar. When the data of several backfilled doses conflict with those of the current dose, which of them anchors the pooled estimate: "highest" (the default here, following Zhao et al., 2024) or "lowest" (the default of be_boin_simulate, following Chen et al., 2026). See the details.
no_slot	Character scalar. What happens to a patient who arrives when neither the dose escalation nor any backfill dose can take a patient: "wait" (the default), the patient waits and is treated as soon as one can, as in tite_boin_simulate; or "leave", the patient is turned away and the trial waits for the next arrival.
accrual	Character scalar. Distribution of the times between arrivals: "exponential" (the default) for a Poisson process, "uniform" for a uniform distribution on (0, 2 / accrual_rate), or "fixed" for arrivals exactly 1 / accrual_rate apart.
dlt_time	Character scalar. Distribution of the time to DLT within the window, and of the time to response within the response window: "weibull" (the default) or "uniform".

late_fraction Numeric scalar strictly between 0 and 1. Share of the DLTs that occur in the second half of the assessment window under the Weibull distribution. Defaults to 0.5.

resp_window Numeric scalar. Length of the window within which a response is observed. Defaults to window.

resp_late_fraction Numeric scalar strictly between 0 and 1. Share of the responses that occur in the second half of the response window under the Weibull distribution. Defaults to 0.5.

resp_cor Numeric scalar strictly between -1 and 1. Correlation of the Gaussian copula that links the DLT and the response of a patient; a positive value makes a response more likely in a patient with a DLT. Defaults to 0, independence.

n_trials Integer scalar. Number of trials to simulate. Defaults to 10000.

start_dose Integer scalar. Dose level for the first cohort. Defaults to 1.

n_earlystop Integer scalar. The trial stops once this many patients, backfilled ones included, have been treated at the current dose and the design would stay there. Defaults to 18.

p_saf Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to 0.6 * target.

p_tox Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to 1.4 * target.

cutoff_eli Numeric scalar. Posterior probability cutoff for dose elimination. Defaults to 0.95.

extrasafe Logical scalar. Apply the stricter safety stopping rule at the lowest dose. Defaults to FALSE.

offset Numeric scalar between 0 and 0.5. Amount by which cutoff_eli is relaxed for the safety stopping rule. Defaults to 0.05.

stay_on_1_of_3 Logical scalar. When TRUE, one DLT out of three patients leads to staying rather than de-escalating, in every use of the decision table. Zhao et al. (2024) use this modification in their main simulation. Defaults to FALSE. See boin_boundary.

bound_mtd Logical scalar. Require the isotonic estimate at the selected dose to be at or below the de-escalation boundary. Defaults to FALSE.

mtd_max_estimate Numeric scalar or NULL. Largest isotonic estimate a dose may have and still be selected as the MTD. See sim_boin.

min_mtd_sample Integer scalar. Smallest number of patients a dose must have received to be eligible as the MTD. Defaults to 1.

overdose_cutoff Numeric scalar or NULL. Doses whose true DLT probability exceeds this value count as overdoses in the overdose component of the result. Defaults to NULL, which uses target.

n_earlystop_rule Character scalar, either "with_stay" or "simple". See boin_simulate.

keep_trials Logical scalar. Keep the trial by trial data of every scenario. Defaults to FALSE.

verbose	Logical scalar. Report progress while the simulations run. Defaults to FALSE.
seed	Integer scalar or NULL. Random seed, applied to every scenario, so that a scenario simulated here matches the same scenario simulated on its own with sim_bf_boin. Defaults to 123.

Value

An object of class c("bf_boin_oc_multi", "backfill_oc_multi", "boin_oc_multi"), which is a list with components

results	Named list of bf_boin_oc objects, one per scenario.
summary_table	Data frame collecting all scenarios, with seven rows per scenario: the true DLT and response rates, the MTD selection, the patients treated, the backfilled patients, the patients with DLT and the average trial duration.
scenario_names	Character vector of scenario names.
n_doses	Number of dose levels.
call	The matched call.

See Also

sim_bf_boin, sim_be_boin_multi

Examples

```
# Scenarios 2 and 6 of Zhao et al. (2024)
scenarios <- list(
  list(name = "Scenario 2", p_true = c(0.12, 0.25, 0.42, 0.49, 0.55),
       p_resp = c(0.20, 0.30, 0.40, 0.50, 0.60)),
  list(name = "Scenario 6", p_true = c(0.12, 0.25, 0.42, 0.49, 0.55),
       p_resp = c(0.30, 0.35, 0.36, 0.36, 0.36))
)

oc <- sim_bf_boin_multi(
  target = 0.25,
  scenarios = scenarios,
  n_cohort = 10,
  cohort_size = 3,
  window = 1,
  accrual_rate = 3,
  n_earlystop = 9,
  stay_on_1_of_3 = TRUE,
  n_trials = 200,
  seed = 123
)
oc
```

`sim_boin` *Operating Characteristics of a BOIN Design*

Description

Simulate a BOIN dose-finding trial many times under one dose-toxicity scenario and summarize how often each dose is selected as the MTD, how many patients are treated at each dose and how many DLTs are observed.

Usage

```
sim_boin(
  target,
  p_true,
  n_cohort,
  cohort_size,
  n_trials = 10000,
  start_dose = 1,
  n_earlystop = 18,
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  titration = FALSE,
  stay_on_1_of_3 = FALSE,
  bound_mtd = FALSE,
  mtd_max_estimate = NULL,
  min_mtd_sample = 1,
  overdose_cutoff = NULL,
  n_earlystop_rule = c("with_stay", "simple"),
  keep_trials = FALSE,
  verbose = FALSE,
  seed = 123
)
```

Arguments

`target`	Numeric scalar. Target DLT probability, for example 0.30.
`p_true`	Numeric vector. True DLT probability at each dose level, in increasing dose order.
`n_cohort`	Integer scalar. Number of cohorts in a trial.
`cohort_size`	Integer scalar or vector. Number of patients per cohort. A scalar is used for every cohort. A vector shorter than `n_cohort` is padded with its last element and a longer one is truncated.
`n_trials`	Integer scalar. Number of trials to simulate. Defaults to 10000.

`start_dose` Integer scalar. Dose level for the first cohort. Defaults to 1. It is ignored when `titration` is `TRUE`, because the titration phase always begins at the lowest dose.

`n_earlystop` Integer scalar. The trial stops once this many patients have been treated at the current dose and the design would stay there. Defaults to 18. Set it to a value above the maximum sample size to switch this rule off.

`p_saf` Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to `0.6 * target`.

`p_tox` Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to `1.4 * target`.

`cutoff_eli` Numeric scalar. Posterior probability cutoff for dose elimination. Defaults to 0.95.

`extrasafe` Logical scalar. Apply the stricter safety stopping rule at the lowest dose. Defaults to `FALSE`.

`offset` Numeric scalar between 0 and 0.5. Amount by which `cutoff_eli` is relaxed for the safety stopping rule. Defaults to 0.05.

`titration` Logical scalar. Start with single patient cohorts until the first DLT is seen. Ignored when the first cohort size is one. Defaults to `FALSE`.

`stay_on_1_of_3` Logical scalar. When `TRUE`, one DLT out of three patients leads to staying at the current dose rather than de-escalating. Defaults to `FALSE`. See `boin_boundary`.

`bound_mtd` Logical scalar. Require the isotonic estimate at the selected dose to be at or below the de-escalation boundary. Defaults to `FALSE`.

`mtd_max_estimate` Numeric scalar or `NULL`. Largest isotonic estimate a dose may have and still be selected as the MTD. Supplying it bounds the selection whatever `bound_mtd` says, and unlike `bound_mtd` it can be set at or below the target rate. It changes only the selection, never the dose-finding itself. Defaults to `NULL`.

`min_mtd_sample` Integer scalar. Smallest number of patients a dose must have received to be eligible as the MTD. Defaults to 1.

`overdose_cutoff` Numeric scalar or `NULL`. Doses whose true DLT probability exceeds this value count as overdoses in the `overdose` component of the result. Defaults to `NULL`, which uses `target`.

`n_earlystop_rule` Character scalar, either `"with_stay"` or `"simple"`. See `boin_simulate`.

`keep_trials` Logical scalar. Keep the full trial by trial data in the result. Defaults to `FALSE`.

`verbose` Logical scalar. Report progress while the simulation runs. Defaults to `FALSE`.

`seed` Integer scalar or `NULL`. Random seed. The state of the calling session is restored on exit. Defaults to 123.

Details

The engine draws one uniform variate per patient, in enrollment order, and applies the decision rules in the order used by `BOIN::get.oc()`. See `boin_simulate` for the two places where a variate

is drawn but not used. With the same seed and matching arguments the two implementations agree trial by trial, not merely on average.

Note that `n_earlystop` defaults to 18 here, whereas the reference implementation defaults to 100, which in practice switches the rule off.

Value

An object of class `boin_oc`, which is a list with components

`sel_percent`	Percentage of trials selecting each dose as the MTD.
`percent_no_mtd`	Percentage of trials ending without an MTD.
`n_pts_dose`	Average number of patients treated at each dose.
`n_tox_dose`	Average number of DLTs observed at each dose.
`total_n_pts`	Average total number of patients per trial.
`total_n_tox`	Average total number of DLTs per trial.
`overdose`	List describing the design's relationship with doses above `overdose_cutoff`. Two questions are answered separately. How far were patients exposed during the trial: `pct_patients` (the percentage of all simulated patients treated at such a dose, that is the probability that a patient is dosed above the cutoff), `pct_patients_by_trial` (the same percentage computed within each trial and then averaged), `avg_n_patients`, `pct_trials_any`, `pct_trials_over_60` and `pct_trials_over_80`. And how often did the design end up recommending such a dose: `pct_trials_mtd_above` (the percentage of all trials whose selected MTD is above the cutoff) and `pct_mtd_above_when_selected` (the same among the trials that selected an MTD at all). The `cutoff` and the dose levels `doses` that exceed it are also returned.
`stop_reason_percent`	Percentage of trials by reason for stopping.
`trials`	Trial level data when `keep_trials` is `TRUE`, otherwise `NULL`.

together with the design parameters and the call.

References

Liu S. and Yuan, Y. (2015). Bayesian Optimal Interval Designs for Phase I Clinical Trials. Journal of the Royal Statistical Society: Series C, 64, 507-523.

See Also

`sim_boin_multi`, `boin_simulate`

Examples

```
oc <- sim_boin(
  target = 0.30,
  p_true = c(0.05, 0.15, 0.30, 0.45, 0.60),
  n_cohort = 10,
  cohort_size = 3,
```

```
  n_trials = 500,
  seed = 123
)
oc


# A larger run with the safety options switched on
oc_safe <- sim_boin(
  target = 0.30,
  p_true = c(0.30, 0.40, 0.50, 0.60, 0.70),
  n_cohort = 20,
  cohort_size = 3,
  n_trials = 10000,
  extrasafe = TRUE,
  bound_mtd = TRUE,
  titration = TRUE,
  seed = 123
)
oc_safe

# How often are patients dosed above a true DLT rate of 0.33?
oc_cut <- sim_boin(
  target = 0.30,
  p_true = c(0.10, 0.20, 0.30, 0.42, 0.55),
  n_cohort = 20,
  cohort_size = 3,
  n_trials = 10000,
  overdose_cutoff = 0.33,
  seed = 123
)
oc_cut$overdose$pct_patients            # patients dosed above 0.33
oc_cut$overdose$pct_trials_mtd_above     # trials recommending a dose above 0.33
```

`sim_boin_multi` *Operating Characteristics Across Several Scenarios*

Description

Run `sim_boin` under a list of dose-toxicity scenarios that share the same design, and collect the results into one table.

Usage

```
sim_boin_multi(
  target,
  scenarios,
  n_cohort,
  cohort_size,
```

```
  n_trials = 10000,
  start_dose = 1,
  n_earlystop = 18,
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  titration = FALSE,
  stay_on_1_of_3 = FALSE,
  bound_mtd = FALSE,
  mtd_max_estimate = NULL,
  min_mtd_sample = 1,
  overdose_cutoff = NULL,
  n_earlystop_rule = c("with_stay", "simple"),
  keep_trials = FALSE,
  verbose = FALSE,
  seed = 123
)
```

Arguments

target	Numeric scalar. Target DLT probability.
scenarios	A list of scenarios. Each element is a list with a `name` and a `p_true` vector, or a plain numeric vector of true DLT probabilities in which case the list names are used as scenario names. Every scenario must have the same number of doses.
n_cohort	Integer scalar. Number of cohorts in a trial.
cohort_size	Integer scalar or vector. Number of patients per cohort.
n_trials	Integer scalar. Number of trials per scenario. Defaults to 10000.
start_dose	Integer scalar. Dose level for the first cohort. Defaults to 1. It is ignored when `titration` is `TRUE`.
n_earlystop	Integer scalar. Early stopping sample size at the current dose. Defaults to 18.
p_saf	Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to `0.6 * target`.
p_tox	Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to `1.4 * target`.
cutoff_eli	Numeric scalar. Posterior probability cutoff for dose elimination. Defaults to 0.95.
extrasafe	Logical scalar. Apply the stricter safety stopping rule at the lowest dose.
offset	Numeric scalar between 0 and 0.5. Relaxation of `cutoff_eli` for the safety stopping rule. Defaults to 0.05.
titration	Logical scalar. Start with single patient cohorts until the first DLT.
stay_on_1_of_3	Logical scalar. Make one DLT out of three patients a stay rather than a de-escalation. See `boin_boundary`.

`bound_mtd` Logical scalar. Bound the isotonic estimate at the selected MTD by the de-escalation boundary.

`mtd_max_estimate` Numeric scalar or `NULL`. Largest isotonic estimate a dose may have and still be selected as the MTD. See `sim_boin`.

`min_mtd_sample` Integer scalar. Smallest number of patients for a dose to be eligible.

`overdose_cutoff` Numeric scalar or `NULL`. Doses whose true DLT probability exceeds this value count as overdoses. Defaults to `NULL`, which uses `target`.

`n_earlystop_rule` Character scalar, either `"with_stay"` or `"simple"`.

`keep_trials` Logical scalar. Keep the trial by trial data of every scenario.

`verbose` Logical scalar. Report progress while the simulations run.

`seed` Integer scalar or `NULL`. Random seed, applied to every scenario, so that a scenario simulated here matches the same scenario simulated on its own with `sim_boin`.

Value

An object of class `boin_oc_multi`, which is a list with components

`results` Named list of `boin_oc` objects, one per scenario.

`summary_table` Data frame collecting all scenarios.

`scenario_names` Character vector of scenario names.

`n_doses` Number of dose levels.

`call` The matched call.

See Also

`sim_boin`

Examples

```
scenarios <- list(
  list(name = "MTD at dose 3", p_true = c(0.05, 0.15, 0.30, 0.45, 0.60)),
  list(name = "All toxic",     p_true = c(0.35, 0.45, 0.55, 0.65, 0.75))
)

oc <- sim_boin_multi(
  target = 0.30,
  scenarios = scenarios,
  n_cohort = 10,
  cohort_size = 3,
  n_trials = 200,
  seed = 123
)
oc
```

`sim_comb_boin` *Operating Characteristics of the BOIN Drug-Combination Design*

Description

Simulate a trial of two drugs given together, run with the BOIN drug-combination design of Lin and Yin (2017), many times under one matrix of true DLT probabilities, and summarize how often each dose combination is selected as the MTD, how many patients are treated at each combination and how many DLTs are observed.

Usage

```
sim_comb_boin(
  target,
  p_true,
  n_cohort,
  cohort_size,
  n_trials = 10000,
  start_dose = c(1, 1),
  n_earlystop = 18,
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  titration = FALSE,
  bound_mtd = FALSE,
  overdose_cutoff = NULL,
  keep_trials = FALSE,
  verbose = FALSE,
  seed = 123
)
```

Arguments

`target`	Numeric scalar. Target DLT probability, for example 0.30.
`p_true`	Numeric matrix. True DLT probability of each dose combination, with one row per dose level of drug A and one column per dose level of drug B, both in increasing dose order.
`n_cohort`	Integer scalar. Number of cohorts in a trial.
`cohort_size`	Integer scalar or vector. Number of patients per cohort. A scalar is used for every cohort. A vector shorter than `n_cohort` is padded with its last element and a longer one is truncated.
`n_trials`	Integer scalar. Number of trials to simulate. Defaults to 10000.
`start_dose`	Integer vector of length two. Dose levels of drug A and drug B for the first cohort. Defaults to `c(1, 1)`.

n_earlystop	Integer scalar. The trial stops once this many patients have been treated at the current combination and the design would neither escalate nor de-escalate, or cannot. Defaults to 18. Set it to a value above the maximum sample size to switch this rule off.
p_saf	Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to `0.6 * target`.
p_tox	Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to `1.4 * target`.
cutoff_eli	Numeric scalar. Posterior probability cutoff for eliminating a combination. Defaults to 0.95.
extrasafe	Logical scalar. Apply the stricter safety stopping rule at the lowest combination. Defaults to `FALSE`.
offset	Numeric scalar between 0 and 0.5. Amount by which `cutoff_eli` is relaxed for the safety stopping rule. Defaults to 0.05.
titration	Logical scalar. Treat one patient per combination, moving up one level of either drug, until the first DLT. Ignored when the first cohort size is one. Defaults to `FALSE`.
bound_mtd	Logical scalar. Require the isotonic estimate at the selected combination to be below the de-escalation boundary. Defaults to `FALSE`.
overdose_cutoff	Numeric scalar or `NULL`. Combinations whose true DLT probability exceeds this value count as overdoses in the `overdose` component of the result. Defaults to `NULL`, which uses `target`.
keep_trials	Logical scalar. Keep the full trial by trial data in the result. Defaults to `FALSE`.
verbose	Logical scalar. Report progress while the simulation runs. Defaults to `FALSE`.
seed	Integer scalar or `NULL`. Random seed. The state of the calling session is restored on exit. Defaults to 123.

Details

The trials are simulated with `comb_boin_simulate` and the MTD is selected with `comb_boin_select_mtd`. A trial stopped for safety at the lowest combination selects none. Both follow `BOIN::get.oc.comb()` with `mtd.contour = FALSE`: with the same seed and matching arguments the two implementations agree trial by trial. `pct_true_mtd_selected` and `pct_patients_true_mtd` are the quantities that the reference implementation reports as the selection percentage of the MTD and the percentage of patients treated at the MTD.

Note that `n_earlystop` defaults to 18 here, whereas the reference implementation defaults to 100, which in practice switches the rule off.

Only the design that searches for a single MTD combination is provided; the waterfall design of Zhang and Yuan (2016), which searches for the MTD contour, is not.

Value

An object of class `comb_boin_oc`, which is a list with components

sel_percent	Matrix of the percentage of trials selecting each combination as the MTD.

`percent_no_mtd` Percentage of trials ending without an MTD.

`n_pts_dose` Matrix of the average number of patients treated at each combination.

`n_tox_dose` Matrix of the average number of DLTs observed at each combination.

`total_n_pts` Average total number of patients per trial.

`total_n_tox` Average total number of DLTs per trial.

`true_mtd` Two-column matrix of the dose levels of drug A and drug B of the combinations whose true DLT probability is closest to `target`.

`pct_true_mtd_selected` Percentage of trials selecting one of these combinations.

`pct_patients_true_mtd` Percentage of all patients treated at one of them.

`overdose` List describing the exposure to combinations above `overdose_cutoff`, with the components of the same list returned by `sim_boin`; `combinations` gives the dose levels of the combinations above the cutoff.

`stop_reason_percent` Percentage of trials by reason for stopping.

`trials` Trial level data when `keep_trials` is `TRUE`, otherwise `NULL`.

together with the design parameters and the call.

References

Lin, R. and Yin, G. (2017). Bayesian Optimal Interval Design for Dose Finding in Drug-Combination Trials. Statistical Methods in Medical Research, 26(5), 2155-2167.

Yan, F., Zhang, L., Zhou, Y., Pan, H., Liu, S. and Yuan, Y. (2020). BOIN: An R Package for Designing Single-Agent and Drug-Combination Dose-Finding Trials Using Bayesian Optimal Interval Designs. Journal of Statistical Software, 94(13), 1-32.

Zhang, L. and Yuan, Y. (2016). A Practical Bayesian Design to Identify the Maximum Tolerated Dose Contour for Drug Combination Trials. Statistics in Medicine, 35(27), 4924-4936.

See Also

`sim_comb_boin_multi`, `comb_boin_simulate`, `comb_boin_select_mtd`

Examples

```
# The example of Yan and colleagues (2020): three levels of drug A, four of
# drug B and a target DLT rate of 0.25
p_true <- matrix(c(0.02, 0.04, 0.08, 0.14,
                   0.08, 0.25, 0.42, 0.48,
                   0.25, 0.45, 0.50, 0.60), nrow = 3, byrow = TRUE)

oc <- sim_comb_boin(
  target = 0.25,
  p_true = p_true,
  n_cohort = 16,
  cohort_size = 3,
```

```
  n_trials = 500,
  seed = 123
)
oc


# The settings of the article, which turn early stopping off
oc_jss <- sim_comb_boin(
  target = 0.25,
  p_true = p_true,
  n_cohort = 16,
  cohort_size = 3,
  n_earlystop = 100,
  n_trials = 1000,
  seed = 6
)
oc_jss$pct_true_mtd_selected
```

`sim_comb_boin_multi` *Operating Characteristics of the BOIN Drug-Combination Design Across Scenarios*

Description

Run `sim_comb_boin` under a list of matrices of true DLT probabilities that share the same design, and collect the main results into one table.

Usage

```
sim_comb_boin_multi(
  target,
  scenarios,
  n_cohort,
  cohort_size,
  n_trials = 10000,
  start_dose = c(1, 1),
  n_earlystop = 18,
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  titration = FALSE,
  bound_mtd = FALSE,
  overdose_cutoff = NULL,
  keep_trials = FALSE,
  verbose = FALSE,
```

```
  seed = 123
)
```

Arguments

target	Numeric scalar. Target DLT probability.
scenarios	A list of scenarios. Each element is a list with a name and a p_true matrix, or a plain matrix of true DLT probabilities in which case the list names are used as scenario names. Every matrix must have the same numbers of rows (levels of drug A) and columns (levels of drug B).
n_cohort	Integer scalar. Number of cohorts in a trial.
cohort_size	Integer scalar or vector. Number of patients per cohort.
n_trials	Integer scalar. Number of trials per scenario. Defaults to 10000.
start_dose	Integer vector of length two. Dose levels of drug A and drug B for the first cohort. Defaults to c(1, 1).
n_earlystop	Integer scalar. Early stopping sample size at the current combination. Defaults to 18.
p_saf	Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to 0.6 * target.
p_tox	Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to 1.4 * target.
cutoff_eli	Numeric scalar. Posterior probability cutoff for eliminating a combination. Defaults to 0.95.
extrasafe	Logical scalar. Apply the stricter safety stopping rule at the lowest combination.
offset	Numeric scalar between 0 and 0.5. Relaxation of cutoff_eli for the safety stopping rule. Defaults to 0.05.
titration	Logical scalar. Treat one patient per combination until the first DLT.
bound_mtd	Logical scalar. Require the isotonic estimate at the selected combination to be below the de-escalation boundary.
overdose_cutoff	Numeric scalar or NULL. Combinations whose true DLT probability exceeds this value count as overdoses. Defaults to NULL, which uses target.
keep_trials	Logical scalar. Keep the trial by trial data of every scenario.
verbose	Logical scalar. Report progress while the simulations run.
seed	Integer scalar or NULL. Random seed, applied to every scenario, so that a scenario simulated here matches the same scenario simulated on its own with sim_comb_boin.

Value

An object of class comb_boin_oc_multi, which is a list with components

results	Named list of comb_boin_oc objects, one per scenario.

summary_table	Data frame with one row per scenario: the true MTD combinations (those closest to target), the percentage of trials selecting one of them, the percentage of patients treated at them, the percentage of trials without an MTD, the percentage of trials selecting a combination above overdose_cutoff, the percentage of patients treated above it, and the average numbers of patients and DLTs.
scenario_names	Character vector of scenario names.
n_rows	Number of dose levels of drug A.
n_cols	Number of dose levels of drug B.
call	The matched call.

See Also

sim_comb_boin

Examples

```
scenarios <- list(
  "Two MTD combinations" = matrix(c(0.05, 0.10, 0.20,
                                    0.10, 0.20, 0.30,
                                    0.20, 0.30, 0.45), 3, byrow = TRUE),
  "All toxic" = matrix(c(0.40, 0.50, 0.60,
                         0.50, 0.60, 0.70,
                         0.60, 0.70, 0.80), 3, byrow = TRUE)
)

oc <- sim_comb_boin_multi(
  target = 0.30,
  scenarios = scenarios,
  n_cohort = 12,
  cohort_size = 3,
  n_trials = 200,
  seed = 123
)
oc

# The full results of one scenario
oc$results[["All toxic"]]
```

---

sim_tite_boin *Operating Characteristics of a TITE-BOIN Design*

Description

Simulate a time-to-event BOIN (TITE-BOIN) trial many times under one dose-toxicity scenario and summarize how often each dose is selected as the MTD, how many patients are treated and how many DLTs are observed at each dose, how long the trials last and how often accrual is suspended.

Usage

```
sim_tite_boin(
  target,
  p_true,
  n_cohort,
  cohort_size,
  window,
  accrual_rate,
  method = c("imputation", "ess"),
  accrual = c("exponential", "uniform", "fixed"),
  dlt_time = c("weibull", "uniform"),
  late_fraction = 0.5,
  prior_weights = c(1, 1, 1)/3,
  max_pending_ratio = NULL,
  min_completed = NULL,
  min_follow_up = 0,
  n_trials = 10000,
  start_dose = 1,
  n_earlystop = 18,
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  bound_mtd = FALSE,
  mtd_max_estimate = NULL,
  min_mtd_sample = 1,
  overdose_cutoff = NULL,
  n_earlystop_rule = c("with_stay", "simple"),
  keep_trials = FALSE,
  verbose = FALSE,
  seed = 123
)
```

Arguments

`target`	Numeric scalar. Target DLT probability, for example 0.30.
`p_true`	Numeric vector. True DLT probability at each dose level, in increasing dose order. The probability of a DLT within the assessment window.
`n_cohort`	Integer scalar. Number of cohorts in a trial.
`cohort_size`	Integer scalar or vector. Number of patients per cohort. A scalar is used for every cohort. A vector shorter than `n_cohort` is padded with its last element and a longer one is truncated.
`window`	Numeric scalar. Length of the DLT assessment window.
`accrual_rate`	Numeric scalar. Average number of patients arriving per unit of time, in the same unit as `window`.

`method` Character scalar, `"imputation"` (the default) or `"ess"`. How the pending patients enter the estimate of the DLT probability. See `tite_boin_decision_table`.

`accrual` Character scalar. Distribution of the times between arrivals: `"exponential"` (the default) for a Poisson process, `"uniform"` for a uniform distribution on `(0, 2 / accrual_rate)`, as in the code of Lin and Yuan (2020), or `"fixed"` for arrivals exactly `1 / accrual_rate` apart.

`dlt_time` Character scalar. Distribution of the time to DLT within the window: `"weibull"` (the default), with the share of DLTs in the second half of the window set by `late_fraction`, or `"uniform"`.

`late_fraction` Numeric scalar strictly between 0 and 1. Share of the DLTs that occur in the second half of the assessment window under the Weibull distribution. Defaults to 0.5, the setting of Yuan et al. (2018) and Lin and Yuan (2020).

`prior_weights` Numeric vector of three non-negative values. Prior probabilities that a DLT occurs in the first, second and last third of the assessment window, used to weight the follow-up of the pending patients in STFT (Yuan et al., 2018, Supplementary Appendix D) and ESS (Lin and Yuan, 2020, Supplementary S1). They are rescaled to sum to one. Defaults to equal weights, which give the unweighted STFT. The decision table does not change with the weights; only the value compared with its boundaries does. This describes the design and is separate from `dlt_time` and `late_fraction`, which describe how the simulated DLTs actually occur.

`max_pending_ratio` Numeric scalar or `NULL`. Suspension rule on the proportion of pending patients. See `tite_boin_decision_table`.

`min_completed` Integer scalar or `NULL`. Number of completed assessments required for escalation. See `tite_boin_decision_table`.

`min_follow_up` Numeric scalar between 0 and 1. Shortest follow-up, as a fraction of `window`, that every pending patient at the current dose must have reached before the dose is escalated; accrual is suspended until then. Defaults to 0, which never suspends on this ground. See `tite_boin_decision_table`.

`n_trials` Integer scalar. Number of trials to simulate. Defaults to 10000.

`start_dose` Integer scalar. Dose level for the first cohort. Defaults to 1.

`n_earlystop` Integer scalar. The trial stops once this many patients have been treated at the current dose and the design would stay there. Defaults to 18.

`p_saf` Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to `0.6 * target`.

`p_tox` Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to `1.4 * target`.

`cutoff_eli` Numeric scalar. Posterior probability cutoff for dose elimination. Defaults to 0.95.

`extrasafe` Logical scalar. Apply the stricter safety stopping rule at the lowest dose. Defaults to `FALSE`.

`offset` Numeric scalar between 0 and 0.5. Amount by which `cutoff_eli` is relaxed for the safety stopping rule. Defaults to 0.05.

bound_mtd — Logical scalar. Require the isotonic estimate at the selected dose to be at or below the de-escalation boundary. Defaults to FALSE.

mtd_max_estimate — Numeric scalar or NULL. Largest isotonic estimate a dose may have and still be selected as the MTD. See sim_boin.

min_mtd_sample — Integer scalar. Smallest number of patients a dose must have received to be eligible as the MTD. Defaults to 1.

overdose_cutoff — Numeric scalar or NULL. Doses whose true DLT probability exceeds this value count as overdoses in the overdose component of the result. Defaults to NULL, which uses target.

n_earlystop_rule — Character scalar, either "with_stay" or "simple". See boin_simulate.

keep_trials — Logical scalar. Keep the full trial by trial data in the result. Defaults to FALSE.

verbose — Logical scalar. Report progress while the simulation runs. Defaults to FALSE.

seed — Integer scalar or NULL. Random seed. When NULL the current state of the random number generator is used. The state of the calling session is restored on exit. Defaults to 123.

Details

The trials are simulated by tite_boin_simulate, whose details describe the conduct of a trial. Once every patient has completed the assessment the MTD is selected from the final data exactly as in sim_boin, with boin_select_mtd, and a trial stopped for safety selects no dose.

When no patient is ever pending at a decision, for example with accrual = "fixed" and window shorter than 1 / accrual_rate, every result agrees with that of sim_boin under the same seed.

Value

An object of class c("tite_boin_oc", "boin_oc"), which has every component of the result of sim_boin, so that code written for sim_boin() also reads it, and in addition

duration_mean — Average trial duration, from the first arrival until every enrolled patient has completed the assessment, in the time unit of window and accrual_rate.

duration_sd — Standard deviation of the trial duration.

pct_trials_suspended — Percentage of trials in which accrual was suspended at least once.

avg_n_suspensions — Average number of suspensions per trial.

avg_time_suspended — Average total time with accrual suspended per trial.

The trials component, when kept, also holds duration, n_suspensions and time_suspended for every trial.

References

Yuan, Y., Lin, R., Li, D., Nie, L. and Warren, K. E. (2018). Time-to-Event Bayesian Optimal Interval Design to Accelerate Phase I Trials. Clinical Cancer Research, 24(20), 4921-4930.

Lin, R. and Yuan, Y. (2020). Time-to-Event Model-Assisted Designs for Dose-Finding Trials with Delayed Toxicity. Biostatistics, 21(4), 807-824.

See Also

tite_boin_simulate, sim_tite_boin_multi, sim_boin

Examples

```
# A three month window with two patients a month
oc <- sim_tite_boin(
  target = 0.30,
  p_true = c(0.05, 0.15, 0.30, 0.45, 0.60),
  n_cohort = 10,
  cohort_size = 3,
  window = 3,
  accrual_rate = 2,
  n_trials = 500,
  seed = 123
)
oc


# The effective sample size of Lin and Yuan (2020)
oc_ess <- sim_tite_boin(
  target = 0.30,
  p_true = c(0.05, 0.15, 0.30, 0.45, 0.60),
  n_cohort = 12,
  cohort_size = 3,
  window = 3,
  accrual_rate = 2,
  method = "ess",
  n_trials = 10000,
  seed = 123
)
oc_ess
```

---

sim_tite_boin_multi *Operating Characteristics of a TITE-BOIN Design Across Several Scenarios*

Description

Run sim_tite_boin under a list of dose-toxicity scenarios that share the same design, and collect the results into one table.

Usage

```
sim_tite_boin_multi(
  target,
  scenarios,
  n_cohort,
  cohort_size,
  window,
  accrual_rate,
  method = c("imputation", "ess"),
  accrual = c("exponential", "uniform", "fixed"),
  dlt_time = c("weibull", "uniform"),
  late_fraction = 0.5,
  prior_weights = c(1, 1, 1)/3,
  max_pending_ratio = NULL,
  min_completed = NULL,
  min_follow_up = 0,
  n_trials = 10000,
  start_dose = 1,
  n_earlystop = 18,
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  bound_mtd = FALSE,
  mtd_max_estimate = NULL,
  min_mtd_sample = 1,
  overdose_cutoff = NULL,
  n_earlystop_rule = c("with_stay", "simple"),
  keep_trials = FALSE,
  verbose = FALSE,
  seed = 123
)
```

Arguments

target	Numeric scalar. Target DLT probability, for example 0.30.
scenarios	A list of scenarios. Each element is a list with a name and a p_true vector, or a plain numeric vector of true DLT probabilities in which case the list names are used as scenario names. Every scenario must have the same number of doses.
n_cohort	Integer scalar. Number of cohorts in a trial.
cohort_size	Integer scalar or vector. Number of patients per cohort. A scalar is used for every cohort. A vector shorter than n_cohort is padded with its last element and a longer one is truncated.
window	Numeric scalar. Length of the DLT assessment window.
accrual_rate	Numeric scalar. Average number of patients arriving per unit of time, in the same unit as window.

`method` Character scalar, `"imputation"` (the default) or `"ess"`. How the pending patients enter the estimate of the DLT probability. See `tite_boin_decision_table`.

`accrual` Character scalar. Distribution of the times between arrivals: `"exponential"` (the default) for a Poisson process, `"uniform"` for a uniform distribution on `(0, 2 / accrual_rate)`, as in the code of Lin and Yuan (2020), or `"fixed"` for arrivals exactly `1 / accrual_rate` apart.

`dlt_time` Character scalar. Distribution of the time to DLT within the window: `"weibull"` (the default), with the share of DLTs in the second half of the window set by `late_fraction`, or `"uniform"`.

`late_fraction` Numeric scalar strictly between 0 and 1. Share of the DLTs that occur in the second half of the assessment window under the Weibull distribution. Defaults to 0.5, the setting of Yuan et al. (2018) and Lin and Yuan (2020).

`prior_weights` Numeric vector of three non-negative values. Prior probabilities that a DLT occurs in the first, second and last third of the assessment window, used to weight the follow-up of the pending patients in STFT (Yuan et al., 2018, Supplementary Appendix D) and ESS (Lin and Yuan, 2020, Supplementary S1). They are rescaled to sum to one. Defaults to equal weights, which give the unweighted STFT. The decision table does not change with the weights; only the value compared with its boundaries does. This describes the design and is separate from `dlt_time` and `late_fraction`, which describe how the simulated DLTs actually occur.

`max_pending_ratio` Numeric scalar or `NULL`. Suspension rule on the proportion of pending patients. See `tite_boin_decision_table`.

`min_completed` Integer scalar or `NULL`. Number of completed assessments required for escalation. See `tite_boin_decision_table`.

`min_follow_up` Numeric scalar between 0 and 1. Shortest follow-up, as a fraction of `window`, that every pending patient at the current dose must have reached before the dose is escalated; accrual is suspended until then. Defaults to 0, which never suspends on this ground. See `tite_boin_decision_table`.

`n_trials` Integer scalar. Number of trials to simulate. Defaults to 10000.

`start_dose` Integer scalar. Dose level for the first cohort. Defaults to 1.

`n_earlystop` Integer scalar. The trial stops once this many patients have been treated at the current dose and the design would stay there. Defaults to 18.

`p_saf` Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to `0.6 * target`.

`p_tox` Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to `1.4 * target`.

`cutoff_eli` Numeric scalar. Posterior probability cutoff for dose elimination. Defaults to 0.95.

`extrasafe` Logical scalar. Apply the stricter safety stopping rule at the lowest dose. Defaults to `FALSE`.

`offset` Numeric scalar between 0 and 0.5. Amount by which `cutoff_eli` is relaxed for the safety stopping rule. Defaults to 0.05.

bound_mtd — Logical scalar. Require the isotonic estimate at the selected dose to be at or below the de-escalation boundary. Defaults to FALSE.

mtd_max_estimate — Numeric scalar or NULL. Largest isotonic estimate a dose may have and still be selected as the MTD. See sim_boin.

min_mtd_sample — Integer scalar. Smallest number of patients a dose must have received to be eligible as the MTD. Defaults to 1.

overdose_cutoff — Numeric scalar or NULL. Doses whose true DLT probability exceeds this value count as overdoses in the overdose component of the result. Defaults to NULL, which uses target.

n_earlystop_rule — Character scalar, either "with_stay" or "simple". See boin_simulate.

keep_trials — Logical scalar. Keep the trial by trial data of every scenario. Defaults to FALSE.

verbose — Logical scalar. Report progress while the simulations run. Defaults to FALSE.

seed — Integer scalar or NULL. Random seed, applied to every scenario, so that a scenario simulated here matches the same scenario simulated on its own with sim_tite_boin. Defaults to 123.

Value

An object of class c("tite_boin_oc_multi", "boin_oc_multi"), which is a list with components

results — Named list of tite_boin_oc objects, one per scenario.

summary_table — Data frame collecting all scenarios, with the four rows of sim_boin_multi and two more per scenario for the average trial duration and the percentage of trials with accrual suspended, both in the last column.

scenario_names — Character vector of scenario names.

n_doses — Number of dose levels.

call — The matched call.

See Also

sim_tite_boin, sim_boin_multi

Examples

```
scenarios <- list(
  list(name = "MTD at dose 3", p_true = c(0.05, 0.15, 0.30, 0.45, 0.60)),
  list(name = "All toxic",     p_true = c(0.35, 0.45, 0.55, 0.65, 0.75))
)

oc <- sim_tite_boin_multi(
  target = 0.30,
  scenarios = scenarios,
  n_cohort = 10,
```

```
  cohort_size = 3,
  window = 3,
  accrual_rate = 2,
  n_trials = 200,
  seed = 123
)
oc
```

`tite_boin_decision_table`

Decision Table for the TITE-BOIN Design

Description

Tabulate the dose assignment rule of the time-to-event Bayesian optimal interval (TITE-BOIN) design, which allows new patients to be treated while the DLT assessment of earlier patients is still pending. For every attainable combination of patients treated, DLTs observed and patients pending at the current dose, the table gives the decision or, when the decision depends on how long the pending patients have been followed, the values of the follow-up statistic at which it changes.

Usage

```
tite_boin_decision_table(
  target,
  max_n,
  method = c("imputation", "ess"),
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  max_pending_ratio = NULL,
  min_completed = NULL,
  min_follow_up = 0
)
```

Arguments

`target`	Numeric scalar. Target DLT probability, for example 0.30.
`max_n`	Integer scalar. Largest number of patients treated at a single dose.
`method`	Character scalar. How the pending patients enter the estimate of the DLT probability: `"imputation"` (the default) for the single mean imputation of Yuan et al. (2018), or `"ess"` for the effective sample size of Lin and Yuan (2020). See the details.
`p_saf`	Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to `0.6 * target`.
`p_tox`	Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to `1.4 * target`.

`cutoff_eli` Numeric scalar. Posterior probability cutoff for dose elimination. Defaults to 0.95.

`max_pending_ratio`
Numeric scalar greater than 0 and at most 1. Accrual is suspended when the proportion of the patients at the current dose whose assessment is pending exceeds this value, unless the dose is de-escalated whatever their outcomes. Defaults to 0.5 for `"imputation"`, following Yuan et al. (2018), and to 1, which never suspends on this ground, for `"ess"`.

`min_completed` Integer scalar. While some patients are pending, escalation requires at least this many patients at the current dose to have completed the assessment, and accrual is suspended instead when fewer have. Defaults to 0 for `"imputation"` and to 2 for `"ess"`, following Lin and Yuan (2020).

`min_follow_up` Numeric scalar between 0 and 1. While some patients are pending, escalation also requires the shortest follow-up among the pending patients at the current dose, as a fraction of the assessment window (MF), to be at least this value, and accrual is suspended instead when it is not. Defaults to 0, which never suspends on this ground. Chen et al. (2025) recommend 0.25, their rule 2 with B = 25. The rule leaves the boundaries on the follow-up statistic unchanged, so it is recorded as an attribute rather than in the rows of the table.

Details

Suppose that n patients have been treated at the current dose, that r of them have completed the DLT assessment with `n_tox` DLTs, and that c = n - r are pending. The standardized total follow-up time (STFT) is the sum of the follow-up times of the pending patients divided by the length of the assessment window, so that STFT lies between 0 and c. With no pending patient the decision is that of the BOIN design, taken from `boin_boundary`.

`method = "imputation"` replaces each pending outcome by its expected value under a uniform time to DLT. The estimate of the DLT probability is `(n_tox + q * (c - STFT)) / n`, where `q = p / (1 - p)` and `p = (n_tox + a) / (r + 1)` is the posterior mean under a Beta(a, 1 - a) prior with `a = target / 2` (Yuan et al., 2018, Supplementary Appendix A). The estimate decreases as STFT grows, so the rule becomes a pair of boundaries on STFT. As in the BOIN design, escalation is possible only while `n_tox / n` is below the target and de-escalation only once it reaches the target.

`method = "ess"` uses the approximated likelihood of Lin and Yuan (2020), under which the estimate is `n_tox / ESS` with the effective sample size ESS = r + STFT, which lies between r and n. The dose is escalated when ESS is at least `n_tox / lambda_e` and de-escalated when it is at most `n_tox / lambda_d`. The boundaries in the table refer to ESS.

The rules are applied in this order: elimination; de-escalation that holds whatever the outcomes of the pending patients, that is when `n_tox / n` is at least `lambda_d`; suspension for too many pending patients; and finally the boundaries on the follow-up statistic, in which an escalation blocked by `min_completed` becomes a suspension. An escalation that remains is then subject to `min_follow_up`. Elimination requires `Pr(p > target | data) > cutoff_eli` computed with all n treated patients, the pending ones counted as without DLT, which is how both articles define it. It is not evaluated before three patients have been treated.

Neither the length of the assessment window nor the weights given to the follow-up times change the table, because both enter only through STFT. A weighted STFT, such as that of a piecewise uniform time to DLT, is compared with the same boundaries.

Value

An object of class `tite_boin_decision_table`, which is a data frame with one row per attainable state of the current dose, ordered by `n`, `n_tox` and `n_pending`, and with the columns

`n`	Number of patients treated at the current dose, from 1 to `max_n`.
`n_tox`	Number of DLTs observed, which can only come from patients who completed the assessment.
`n_pending`	Number of patients whose assessment is pending.
`decision`	Decision code. `"E"` escalates, `"S"` stays, `"D"` de-escalates, `"DE"` de-escalates and eliminates the dose together with all higher doses, and `"SUS"` suspends accrual until more data are available. These do not depend on the follow-up of the pending patients. The codes `"E/S"`, `"S/D"` and `"E/S/D"` mark states whose decision depends on the follow-up statistic through `esc_bound` and `deesc_bound`; in them `"SUS"` takes the place of `"E"` when escalation is blocked by `min_completed`.
`esc_bound`	Escalate, or suspend when escalation is blocked, when the follow-up statistic is at least this value. `NA` when no such boundary applies.
`deesc_bound`	De-escalate when the follow-up statistic is at most this value. `NA` when no such boundary applies.

The design parameters are stored as attributes: `method`, `statistic` (`"STFT"` or `"ESS"`), `target`, `p_saf`, `p_tox`, `lambda_e`, `lambda_d`, `cutoff_eli`, `max_pending_ratio`, `min_completed` and `min_follow_up`. When `min_follow_up` is positive, an escalation (code `"E"`, or `"E"` as the first part of a code with a slash) in a state with pending patients takes place only when MF is at least `min_follow_up`, and accrual is suspended otherwise. The object has `print` and `plot` methods.

Agreement with the published tables

With the defaults, `method = "imputation"` reproduces every entry of Table 1 (target 0.2) of Yuan et al. (2018) and of Table S1 (target 0.3) of its supplementary appendix. These tables allow de-escalation when `n_tox / n` equals the target (Table 1, 15 patients with 3 DLTs), which equation (5) of the appendix, read literally, would not. The implementation follows the tables, which agree with the verbal description of the rule in the appendix.

With `max_pending_ratio = 0.49` and `min_follow_up = 0.25`, `method = "imputation"` reproduces every entry of Table 1 (target 0.3, up to six patients) of Chen et al. (2025) and of Table A1 (target 0.25, up to nine patients) of the supplementary materials of Chen et al. (2026). Their rule 1 suspends accrual when fewer than 51 percent of the patients have completed the assessment, which for whole numbers of patients is the same as more than 49 percent pending, unless the observed DLT rate already exceeds `lambda_d`; this exception is the de-escalation that holds whatever the outcomes of the pending patients.

For `method = "ess"` no table for the BOIN design has been published; the tables of Lin and Yuan (2020) are for the keyboard and mTPI designs.

References

Yuan, Y., Lin, R., Li, D., Nie, L. and Warren, K. E. (2018). Time-to-Event Bayesian Optimal Interval Design to Accelerate Phase I Trials. Clinical Cancer Research, 24(20), 4921-4930.

Lin, R. and Yuan, Y. (2020). Time-to-Event Model-Assisted Designs for Dose-Finding Trials with Delayed Toxicity. Biostatistics, 21(4), 807-824.

Chen, K., Chen, T.-Y., Zhang, Y., Lin, R. and Yuan, Y. (2025). Practical Considerations for Using the TITE-BOIN Design to Handle Late-Onset Toxicity or Fast Accrual in Phase I Trials. Clinical Cancer Research, 31(13), 2573-2580.

Chen, K., Zhao, Y., Takeda, K. and Yuan, Y. (2026). BE-BOIN: A Dose Optimization Design Accommodating Backfill and Late-Onset Toxicity. Therapeutic Innovation and Regulatory Science. doi:10.1007/s43441026009940

See Also

print.tite_boin_decision_table, plot.tite_boin_decision_table, boin_decision_table

Examples

```
# Table 1 of Yuan et al. (2018)
decisions <- tite_boin_decision_table(target = 0.2, max_n = 15)
print(decisions, cohort_size = 3)

# Nine patients, one DLT and four pending: escalate once STFT reaches 2.15
decisions[decisions$n == 9 & decisions$n_tox == 1 & decisions$n_pending == 4, ]

# Table 1 of Chen et al. (2025): at least 51 percent of the patients with a
# completed assessment, and escalation only once every pending patient has
# been followed for a quarter of the window
decisions_chen <- tite_boin_decision_table(target = 0.3, max_n = 6,
                                           max_pending_ratio = 0.49,
                                           min_follow_up = 0.25)
print(decisions_chen, cohort_size = 3)

# The effective sample size of Lin and Yuan (2020)
decisions_ess <- tite_boin_decision_table(target = 0.3, max_n = 12, method = "ess")
print(decisions_ess, cohort_size = 3)
```

---

tite_boin_simulate *Simulate TITE-BOIN Trials*

Description

Run the time-to-event BOIN (TITE-BOIN) design over many simulated trials, in which patients arrive over time and new patients are treated while the DLT assessment of earlier patients is still pending. Return the patient and DLT counts at every dose together with the duration of each trial and the time spent with accrual suspended.

Usage

```
tite_boin_simulate(
  target,
  p_true,
  n_cohort,
  cohort_size,
  window,
  accrual_rate,
  method = c("imputation", "ess"),
  accrual = c("exponential", "uniform", "fixed"),
  dlt_time = c("weibull", "uniform"),
  late_fraction = 0.5,
  prior_weights = c(1, 1, 1)/3,
  max_pending_ratio = NULL,
  min_completed = NULL,
  min_follow_up = 0,
  n_trials = 10000,
  start_dose = 1,
  n_earlystop = 18,
  p_saf = NULL,
  p_tox = NULL,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  n_earlystop_rule = c("with_stay", "simple"),
  seed = 123
)
```

Arguments

`target`	Numeric scalar. Target DLT probability, for example 0.30.
`p_true`	Numeric vector. True DLT probability at each dose level, in increasing dose order. The probability of a DLT within the assessment window.
`n_cohort`	Integer scalar. Number of cohorts in a trial.
`cohort_size`	Integer scalar or vector. Number of patients per cohort. A scalar is used for every cohort. A vector shorter than `n_cohort` is padded with its last element and a longer one is truncated.
`window`	Numeric scalar. Length of the DLT assessment window.
`accrual_rate`	Numeric scalar. Average number of patients arriving per unit of time, in the same unit as `window`.
`method`	Character scalar, `"imputation"` (the default) or `"ess"`. How the pending patients enter the estimate of the DLT probability. See `tite_boin_decision_table`.
`accrual`	Character scalar. Distribution of the times between arrivals: `"exponential"` (the default) for a Poisson process, `"uniform"` for a uniform distribution on `(0, 2 / accrual_rate)`, as in the code of Lin and Yuan (2020), or `"fixed"` for arrivals exactly `1 / accrual_rate` apart.

`dlt_time` Character scalar. Distribution of the time to DLT within the window: `"weibull"` (the default), with the share of DLTs in the second half of the window set by `late_fraction`, or `"uniform"`.

`late_fraction` Numeric scalar strictly between 0 and 1. Share of the DLTs that occur in the second half of the assessment window under the Weibull distribution. Defaults to 0.5, the setting of Yuan et al. (2018) and Lin and Yuan (2020).

`prior_weights` Numeric vector of three non-negative values. Prior probabilities that a DLT occurs in the first, second and last third of the assessment window, used to weight the follow-up of the pending patients in STFT (Yuan et al., 2018, Supplementary Appendix D) and ESS (Lin and Yuan, 2020, Supplementary S1). They are rescaled to sum to one. Defaults to equal weights, which give the unweighted STFT. The decision table does not change with the weights; only the value compared with its boundaries does. This describes the design and is separate from `dlt_time` and `late_fraction`, which describe how the simulated DLTs actually occur.

`max_pending_ratio` Numeric scalar or `NULL`. Suspension rule on the proportion of pending patients. See `tite_boin_decision_table`.

`min_completed` Integer scalar or `NULL`. Number of completed assessments required for escalation. See `tite_boin_decision_table`.

`min_follow_up` Numeric scalar between 0 and 1. Shortest follow-up, as a fraction of `window`, that every pending patient at the current dose must have reached before the dose is escalated; accrual is suspended until then. Defaults to 0, which never suspends on this ground. See `tite_boin_decision_table`.

`n_trials` Integer scalar. Number of trials to simulate. Defaults to 10000.

`start_dose` Integer scalar. Dose level for the first cohort. Defaults to 1.

`n_earlystop` Integer scalar. The trial stops once this many patients have been treated at the current dose and the design would stay there. Defaults to 18.

`p_saf` Numeric scalar. Highest DLT probability deemed subtherapeutic. Defaults to `0.6 * target`.

`p_tox` Numeric scalar. Lowest DLT probability deemed overly toxic. Defaults to `1.4 * target`.

`cutoff_eli` Numeric scalar. Posterior probability cutoff for dose elimination. Defaults to 0.95.

`extrasafe` Logical scalar. Apply the stricter safety stopping rule at the lowest dose. Defaults to `FALSE`.

`offset` Numeric scalar between 0 and 0.5. Amount by which `cutoff_eli` is relaxed for the safety stopping rule. Defaults to 0.05.

`n_earlystop_rule` Character scalar, either `"with_stay"` or `"simple"`. See `boin_simulate`.

`seed` Integer scalar or `NULL`. Random seed. When `NULL` the current state of the random number generator is used. The state of the calling session is restored on exit. Defaults to 123.

Details

Patients arrive one at a time. The dose for a cohort is decided when its first patient arrives, from the data available at that moment, and the other patients of the cohort receive the same dose. The decision follows `tite_boin_decision_table`. A patient with a DLT completes the assessment when the DLT occurs, and a patient without one at the end of the window. When the decision is to suspend accrual, the arriving patient waits until the next pending patient at the current dose completes the assessment, or, with a positive `min_follow_up`, until the last pending patient to arrive at the current dose has been followed for that fraction of the window, whichever comes first, and the decision is taken again at that moment. Later patients arrive after the one who waited.

At every decision the rules are applied in the order used by `boin_simulate`: elimination at the current dose and the extra safety rule, early stopping on `n_earlystop`, then the dose transition. Because DLTs can be observed after a dose has been left, the elimination rule is also checked at the other doses; a dose that meets it is eliminated together with all higher doses, and the trial moves below it. Elimination counts every treated patient, the pending ones as without DLT.

One uniform variate is drawn from R's random number stream per patient, in enrollment order and in the cohort blocks of `boin_simulate`. It decides whether the patient has a DLT and, under the Weibull distribution, gives the time to DLT through the inverse distribution function, which falls inside the window exactly when there is a DLT. The times between arrivals come from a second generator (xoshiro256**) seeded with `seed`, or with one draw from R's stream when `seed` is `NULL`. As a result, whenever no patient is pending at any decision, for example with `accrual = "fixed"` and `window` shorter than `1 / accrual_rate`, the trials are identical to those of `boin_simulate` with the same seed.

The follow-up of a pending patient counts towards STFT as the share of the window already covered, or, with `prior_weights`, as the prior probability that a DLT would have occurred by then. Both articles evaluate their designs with equal weights, the default.

The titration phase and the `stay_on_1_of_3` option of `boin_simulate` are not available, as neither article defines them for the time-to-event design.

Value

An object of class `tite_boin_trials`, which is a list with components

`n_pts`	Integer matrix, `n_trials` by number of doses, of patients treated.
`n_tox`	Integer matrix of the same shape, of DLTs, counted once every patient has completed the assessment.
`eliminated`	Logical matrix of the same shape, marking doses eliminated during the trial.
`cohorts_used`	Integer vector of cohorts treated in each trial.
`stop_reason`	Character vector giving why each trial ended, with the values of `boin_simulate`.
`duration`	Numeric vector. Time from the first arrival until every enrolled patient has completed the assessment.
`n_suspensions`	Integer vector. Number of decisions at which accrual had to be suspended.
`time_suspended`	Numeric vector. Total time with accrual suspended.
`boundary`	The `boin_boundary` object used.
`settings`	List of the design parameters.

References

Yuan, Y., Lin, R., Li, D., Nie, L. and Warren, K. E. (2018). Time-to-Event Bayesian Optimal Interval Design to Accelerate Phase I Trials. Clinical Cancer Research, 24(20), 4921-4930.

Lin, R. and Yuan, Y. (2020). Time-to-Event Model-Assisted Designs for Dose-Finding Trials with Delayed Toxicity. Biostatistics, 21(4), 807-824.

See Also

tite_boin_decision_table, boin_simulate

Examples

```
# A three month window with two patients a month
trials <- tite_boin_simulate(
  target = 0.30,
  p_true = c(0.05, 0.15, 0.30, 0.45, 0.60),
  n_cohort = 10,
  cohort_size = 3,
  window = 3,
  accrual_rate = 2,
  n_trials = 200,
  seed = 123
)
trials

summary(trials$duration)
```

Index

as.data.frame.boin_boundary, 3

be_boin_simulate, 4, *9*, *12*, *43, 44*, *63*, *69*, *73*
bf_boin_simulate, *5*, *7*, 8, *43, 44*, *61*, *65*, *70, 71*
boin_boundary, *3*, *10*, 13, *15, 16*, *19*, *23*, *25, 26*, *31*, *44*, *70*, *74*, *77*, *80*, *96*
boin_decision_table, *14*, 15, *27*, *34*, *36, 37*, *45*, *98*
boin_isotonic, 16, *21*
boin_lambda, *14*, 18, *19*
boin_p_tox, 19
boin_select_mtd, *17*, 20, *24*, *90*
boin_simulate, *6*, *10–12*, *21*, 22, *48*, *63*, *66*, *70*, *74*, *77, 78*, *90*, *94*, *100–102*
boin_stopping_table, 24

comb_boin_next, 25, *37, 38*, *49*
comb_boin_select_mtd, 28, *32*, *83, 84*
comb_boin_simulate, *25*, *28*, 30, *30*, *51*, *83, 84*

decision_table_3p3, 33, *38, 39*, *52, 53*

oc_3p3, *33*, 34, *34*, *53, 54*, *58, 59*

plot.boin_decision_table, *16*, 36, *37*, *39, 40*, *45*
plot.comb_boin_next, *27*, 37
plot.decision_table_3p3, *34*, 38, *53*
plot.tite_boin_decision_table, 39, *55*, *98*
print.backfill_oc, 41
print.backfill_oc_multi, 42
print.backfill_trials, 43
print.boin_boundary, 44
print.boin_decision_table, *37*, 45
print.boin_oc, 46, *53*, *55*
print.boin_oc_multi, 47
print.boin_trials, 48
print.comb_boin_next, 49
print.comb_boin_oc, 49
print.comb_boin_oc_multi, 50
print.comb_boin_trials, 51
print.decision_table_3p3, *34*, *39*, 52
print.oc_3p3, 53
print.tite_boin_decision_table, *41*, 54, *98*
print.tite_boin_oc, 55
print.tite_boin_oc_multi, 56
print.tite_boin_trials, 58

sim_3p3, *33*, *36*, *53, 54*, 58
sim_be_boin, *41, 42*, 60, *64*, *67*, *71*
sim_be_boin_multi, *42, 43*, *63*, 64, *75*
sim_bf_boin, *41, 42*, *63*, 67, *72*, *75*
sim_bf_boin_multi, *42, 43*, *67*, *71*, 72
sim_boin, *22*, *24*, *34–36*, *46*, *59*, *62*, *66*, *70, 71*, *74*, 76, *79*, *81*, *84*, *90, 91*, *94*
sim_boin_multi, *47*, *78*, 79, *94*
sim_comb_boin, *25*, *27*, *30*, *32*, *49*, 82, *85–87*
sim_comb_boin_multi, *50*, *84*, 85
sim_tite_boin, *55*, *63*, 87, *91*, *94*
sim_tite_boin_multi, *56*, 91, *91*

tite_boin_decision_table, *6, 7*, *39*, *41*, *54, 55*, *62*, *66*, *89*, *93*, 95, *99–102*
tite_boin_simulate, *5*, *7*, *9*, *11*, *58*, *61*, *65*, *69*, *73*, *90, 91*, 98