

Package ‘tolerance’

September 3, 2014

Type Package

Title Functions for calculating tolerance intervals

Version 1.0.0

Date 2014-09-03

Depends rgl, stats4

Description Tolerance limits provide the limits between which we can expect to find a specified proportion of a sampled population with a given level of confidence. This package provides functions for estimating tolerance limits for various distributions. Plotting is also available for tolerance limits of continuous random variables.

License GPL (>= 2)

Author Derek S. Young [aut, cre]

Maintainer Derek S. Young <derek.young@uky.edu>

NeedsCompilation no

Repository CRAN

Date/Publication 2014-09-03 15:55:50

R topics documented:

tolerance-package	2
acc.samp	3
anovatul.int	5
bintol.int	7
bonftol.int	9
cautol.int	11
diffnormtol.int	12
DiffProp	14
distfree.est	15
exp2tol.int	17

exptol.int	18
exttol.int	20
F1	22
fidbintol.int	23
fidnegbintol.int	25
fidpoistol.int	27
gamtol.int	29
hypertol.int	31
K.factor	32
K.table	34
laptol.int	36
logistol.int	37
mvregtol.region	38
mvtol.region	40
negbintol.int	42
NegHypergeometric	44
neghypertol.int	46
nlregtol.int	47
normtol.int	49
np.order	51
npregtol.int	52
nptol.int	54
paretotol.int	56
plottol	58
poistol.int	60
regtol.int	62
TwoParExponential	64
umatol.int	65
uniftol.int	66
ZipfMandelbrot	68
zipftol.int	70
zm.ll	72

Index	75
--------------	-----------

tolerance-package	<i>Functions for Calculating Tolerance Intervals</i>
-------------------	--

Description

A collection of functions for calculating $[100(1-\alpha)\%, 100(P)\%]$ tolerance intervals, which are intervals with $100(1-\alpha)\%$ confidence of covering $100(P)\%$ of the sampled population of interest.

Details

Package: tolerance
Type: Package
Version: 0.1.0
Date: 2014-09-03
Depends: rgl, stats4
License: GPL (>= 2)

Author(s)

Derek S. Young, Ph.D.

Maintainer: Derek S. Young <derek.young@uky.edu>

References

Hahn, G. J. and Meeker, W. Q. (1991), *Statistical Intervals: A Guide for Practitioners*, Wiley-Interscience.

Krishnamoorthy, K. and Mathew, T. (2009), *Statistical Tolerance Regions: Theory, Applications, and Computation*, Wiley.

Patel, J. K. (1986), Tolerance Intervals - A Review, *Communications in Statistics - Theory and Methodology*, **15**, 2719–2762.

Young, D. S. (2010), tolerance: An R Package for Estimating Tolerance Intervals, *Journal of Statistical Software*, **36**(5), 1–39.

See Also

[confint](#)

acc.samp

Acceptance Sampling

Description

Provides an upper bound on the number of acceptable rejects or nonconformities in a process. This is similar to a 1-sided upper tolerance bound for a hypergeometric random variable.

Usage

```
acc.samp(n, N, alpha = 0.05, P = 0.99, AQL = 0.01, RQL = 0.02)
```

Arguments

n	The sample size to be drawn from the inventory.
N	The total inventory (or lot) size.
alpha	1-alpha is the confidence level for bounding the probability of accepting the inventory.
P	The proportion of items in the inventory which are to be accountable.
AQL	The acceptable quality level, which is the largest proportion of defects in a process considered acceptable. Note that $0 < \text{AQL} < 1$.
RQL	The rejectable quality level, which is the largest proportion of defects in an independent lot that one is willing to tolerate. Note that $\text{AQL} < \text{RQL} < 1$.

Value

acc.samp returns a matrix with the following quantities:

acceptance.limit	The number of items in the sample which may be unaccountable, yet still be able to attain the desired confidence level 1-alpha.
lot.size	The total inventory (or lot) size N.
confidence	The confidence level 1-alpha.
P	The proportion of accountable items specified by the user.
AQL	The acceptable quality level as specified by the user. If the sampling were to be repeated numerous times as a process, then this quantity specifies the proportion of missing items considered acceptable from the process as a whole. Conditioning on the calculated value for acceptance.limit, the AQL is used to estimate the producer's risk (see prod.risk below).
RQL	The rejectable quality level as specified by the user. This is the proportion of individual items in a sample one is willing to tolerate missing. Conditioning on the calculated value for acceptance.limit, the RQL is used to estimate the consumer's risk (see cons.risk below).
sample.size	The sample size drawn as specified by n.
prod.risk	The producer's risk at the specified AQL. This is the probability of rejecting an audit of a good inventory (also called the Type I error). A good inventory can be rejected if an unfortunate random sample is selected (e.g., most of the missing items happened to be selected for the audit). 1-prod.risk gives the confidence level of this sampling plan for the specified AQL and RQL. If it is lower than the confidence level desired (e.g., because the AQL is too high), then a warning message will be displayed.
cons.risk	The consumer's risk at the specified RQL. This is the probability of accepting an audit of a bad inventory (also called the Type II error). A bad inventory can be accepted if a fortunate random sample is selected (e.g., most of the missing items happened to not be selected for the audit).

References

Montgomery, D. C. (2005), *Introduction to Statistical Quality Control*, Fifth Edition, John Wiley & Sons, Inc.

See Also

[Hypergeometric](#)

Examples

```
## A 90%/90% acceptance sampling plan for a sample of 450
## drawn from a lot size of 960.

acc.samp(n = 450, N = 960, alpha = 0.10, P = 0.90, AQL = 0.07,
        RQL = 0.10)
```

 anovatol.int

Tolerance Intervals for ANOVA

Description

Tolerance intervals for each factor level in a balanced (or nearly-balanced) ANOVA.

Usage

```
anovatol.int(lm.out, data, alpha = 0.05, P = 0.99, side = 1,
            method = c("HE", "HE2", "WBE", "ELL", "KM",
                    "EXACT", "OCT"), m = 50)
```

Arguments

lm.out	An object of class <code>lm</code> (i.e., the results from the linear model fitting routine such that the <code>anova</code> function can act upon).
data	A data frame consisting of the data fitted in <code>lm.out</code> . Note that <code>data</code> must have one column for each main effect (i.e., factor) that is analyzed in <code>lm.out</code> and that these columns must be of class <code>factor</code> .
alpha	The level chosen such that $1-\alpha$ is the confidence level.
P	The proportion of the population to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by <code>side = 1</code> or <code>side = 2</code> , respectively).

method	The method for calculating the k-factors. The k-factor for the 1-sided tolerance intervals is performed exactly and thus is the same for the chosen method. "HE" is the Howe method and is often viewed as being extremely accurate, even for small sample sizes. "HE2" is a second method due to Howe, which performs similarly to the Weissberg-Beatty method, but is computationally simpler. "WBE" is the Weissberg-Beatty method (also called the Wald-Wolfowitz method), which performs similarly to the first Howe method for larger sample sizes. "ELL" is the Ellison correction to the Weissberg-Beatty method when f is appreciably larger than n^2 . A warning message is displayed if f is not larger than n^2 . "KM" is the Krishnamoorthy-Mathew approximation to the exact solution, which works well for larger sample sizes. "EXACT" computes the k-factor exactly by finding the integral solution to the problem via the <code>integrate</code> function. Note the computation time of this method is largely determined by m . "OCT" is the Owen approach to compute the k-factor when controlling the tails so that there is not more than $(1-P)/2$ of the data in each tail of the distribution.
m	The maximum number of subintervals to be used in the <code>integrate</code> function. This is necessary only for <code>method = "EXACT"</code> and <code>method = "OCT"</code> . The larger the number, the more accurate the solution. Too low of a value can result in an error. A large value can also cause the function to be slow for <code>method = "EXACT"</code> .

Value

`anovatol.int` returns a list where each element is a data frame corresponding to each main effect (i.e., factor) tested in the ANOVA and the rows of each data frame are the levels of that factor. The columns of each data frame report the following:

mean	The mean for that factor level.
n	The effective sample size for that factor level.
k	The k-factor for constructing the respective factor level's tolerance interval.
1-sided.lower	The 1-sided lower tolerance bound. This is given only if <code>side = 1</code> .
1-sided.upper	The 1-sided upper tolerance bound. This is given only if <code>side = 1</code> .
2-sided.lower	The 2-sided lower tolerance bound. This is given only if <code>side = 2</code> .
2-sided.upper	The 2-sided upper tolerance bound. This is given only if <code>side = 2</code> .

References

- Howe, W. G. (1969), Two-Sided Tolerance Limits for Normal Populations - Some Improvements, *Journal of the American Statistical Association*, **64**, 610–620.
- Weissberg, A. and Beatty, G. (1969), Tables of Tolerance Limit Factors for Normal Distributions, *Technometrics*, **2**, 483–500.

See Also

[K.factor](#), [normtol.int](#), [lm](#), [anova](#)

Examples

```
## 90%/95% 2-sided tolerance intervals for a 2-way ANOVA
## using the "warpbreaks" data.

attach(warpbreaks)

lm.out <- lm(breaks ~ wool + tension)
out <- anovitol.int(lm.out, data = warpbreaks, alpha = 0.10,
                   P = 0.95, side = 2, method = "HE")

out

plottol(out, x = warpbreaks)
```

bintol.int

Binomial Tolerance Intervals

Description

Provides 1-sided or 2-sided tolerance intervals for binomial random variables. From a statistical quality control perspective, these limits use the proportion of defective (or acceptable) items in a sample to bound the number of defective (or acceptable) items in future productions of a specified quantity.

Usage

```
bintol.int(x, n, m = NULL, alpha = 0.05, P = 0.99, side = 1,
           method = c("LS", "WS", "AC", "JF", "CP", "AS",
                     "LO", "PR", "PO", "CL", "CC", "CWS"),
           a1 = 0.5, a2 = 0.5)
```

Arguments

x	The number of defective (or acceptable) units in the sample. Can be a vector of length n, in which case the sum of x is used.
n	The size of the random sample of units selected for inspection.
m	The quantity produced in future groups. If m = NULL, then the tolerance limits will be constructed assuming n for this quantity.
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of the defective (or acceptable) units in future samples of size m to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by side = 1 or side = 2, respectively).

method	The method for calculating the lower and upper confidence bounds, which are used in the calculation of the tolerance bounds. The default method is "LS", which is the large-sample method. "WS" is Wilson's method, which is just the score confidence interval. "AC" gives the Agresti-Coull method, which is also appropriate when the sample size is large. "JF" is Jeffreys' method, which is a Bayesian approach to the estimation. "CP" is the Clopper-Pearson (exact) method, which is based on beta percentiles and provides a more conservative interval. "AS" is the arcsine method, which is appropriate when the sample proportion is not too close to 0 or 1. "LO" is the logit method, which also is appropriate when the sample proportion is not too close to 0 or 1, but yields a more conservative interval. "PR" uses a probit transformation and is accurate for large sample sizes. "PO" is based on a Poisson parameterization, but it tends to be more erratic compared to the other methods. "CL" is the complementary log transformation and also tends to perform well for large sample sizes. "CC" gives a continuity-corrected version of the large-sample method. "CWS" gives a continuity-corrected version of Wilson's method. More information on these methods can be found in the "References".
a1	This specifies the first shape hyperparameter when using Jeffreys' method.
a2	This specifies the second shape hyperparameter when using Jeffreys' method.

Value

bintol.int returns a data frame with items:

alpha	The specified significance level.
P	The proportion of defective (or acceptable) units in future samples of size m .
p.hat	The proportion of defective (or acceptable) units in the sample, calculated by x/n .
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

References

Brown, L. D., Cai, T. T., and DasGupta, A. (2001), Interval Estimation for a Binomial Proportion, *Statistical Science*, **16**, 101–133.

Hahn, G. J. and Chandra, R. (1981), Tolerance Intervals for Poisson and Binomial Variables, *Journal of Quality Technology*, **13**, 100–110.

Newcombe, R. G. (1998), Two-Sided Confidence Intervals for the Single Proportion: Comparison of Seven Methods, *Statistics in Medicine*, **17**, 857–872.

See Also

[Binomial](#), [umatol.int](#)

Examples

```

## 85%/90% 2-sided binomial tolerance intervals for a future
## lot of 2500 when a sample of 230 were drawn from a lot of
## 1000. All methods but Jeffreys' method are compared
## below.

bintol.int(x = 230, n = 1000, m = 2500, alpha = 0.15, P = 0.90,
           side = 2, method = "LS")
bintol.int(x = 230, n = 1000, m = 2500, alpha = 0.15, P = 0.90,
           side = 2, method = "WS")
bintol.int(x = 230, n = 1000, m = 2500, alpha = 0.15, P = 0.90,
           side = 2, method = "AC")
bintol.int(x = 230, n = 1000, m = 2500, alpha = 0.15, P = 0.90,
           side = 2, method = "CP")
bintol.int(x = 230, n = 1000, m = 2500, alpha = 0.15, P = 0.90,
           side = 2, method = "AS")
bintol.int(x = 230, n = 1000, m = 2500, alpha = 0.15, P = 0.90,
           side = 2, method = "LO")
bintol.int(x = 230, n = 1000, m = 2500, alpha = 0.15, P = 0.90,
           side = 2, method = "PR")
bintol.int(x = 230, n = 1000, m = 2500, alpha = 0.15, P = 0.90,
           side = 2, method = "PO")
bintol.int(x = 230, n = 1000, m = 2500, alpha = 0.15, P = 0.90,
           side = 2, method = "CL")
bintol.int(x = 230, n = 1000, m = 2500, alpha = 0.15, P = 0.90,
           side = 2, method = "CC")
bintol.int(x = 230, n = 1000, m = 2500, alpha = 0.15, P = 0.90,
           side = 2, method = "CWS")

## Using Jeffreys' method to construct the 85%/90% 1-sided
## binomial tolerance limits. The first calculation assumes
## a prior on the proportion of defects which places greater
## density on values near 0. The second calculation assumes
## a prior on the proportion of defects which places greater
## density on values near 1.

bintol.int(x = 230, n = 1000, m = 2500, alpha = 0.15, P = 0.90,
           side = 1, method = "JF", a1 = 2, a2 = 10)
bintol.int(x = 230, n = 1000, m = 2500, alpha = 0.15, P = 0.90,
           side = 1, method = "JF", a1 = 5, a2 = 1)

```

Description

This function allows the user to control what proportion of the population is to be in the tails of the given distribution for a 2-sided tolerance interval. The result is a conservative approximation based on Bonferroni's inequality.

Usage

```
bonftol.int(fn, P1 = 0.005, P2 = 0.005, alpha = 0.05, ...)
```

Arguments

fn	The function name for the 2-sided tolerance interval to be calculated.
P1	The proportion of the population not covered in the lower tail of the distribution.
P2	The proportion of the population not covered in the upper tail of the distribution.
alpha	The level chosen such that $1 - \alpha$ is the confidence level.
...	Additional arguments passed to fn, including the data. All arguments that would be specified in fn must also be specified here.

Value

The results for the 2-sided tolerance interval procedure are reported. See the corresponding help file for fn about specific output. Note that the (minimum) proportion of the population to be covered by this interval is $1 - (P1 + P2)$.

Note

This function can be used with any 2-sided tolerance interval function, including the regression tolerance interval functions.

References

Jensen, W. A. (2009), Approximations of Tolerance Intervals for Normally Distributed Data, *Quality and Reliability Engineering International*, **25**, 571–580.

Patel, J. K. (1986), Tolerance Intervals - A Review, *Communications in Statistics - Theory and Methodology*, **15**, 2719–2762.

Examples

```
## 95%/97% tolerance interval for normally distributed
## data controlling 1% of the data is in the lower tail
## and 2% of the data in the upper tail.

set.seed(100)
x <- rnorm(100, 0, 0.2)
bonftol.int(normtol.int, x = x, P1 = 0.01, P2 = 0.02,
            alpha = 0.05, method = "HE")
```

cautol.int	<i>Cauchy Tolerance Intervals</i>
------------	-----------------------------------

Description

Provides 1-sided or 2-sided tolerance intervals for Cauchy distributed data.

Usage

```
cautol.int(x, alpha = 0.05, P = 0.99, side = 1)
```

Arguments

x	A vector of data which is Cauchy distributed.
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of the population to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by side = 1 or side = 2, respectively).

Value

cautol.int returns a data.frame with items:

alpha	The specified significance level.
P	The proportion of the population covered by this tolerance interval.
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

References

Bain, L. J. (1978), *Statistical Analysis of Reliability and Life-Testing Models*, Marcel Dekker, Inc.

See Also

[Cauchy](#)

Examples

```
## 95%/90% 2-sided Cauchy tolerance interval for a sample
## of size 1000.

set.seed(100)
x <- rcauchy(1000, 100000, 10)
out <- cauto1.int(x = x, alpha = 0.05, P = 0.90, side = 2)
out

plottol(out, x, plot.type = "both", x.lab = "Cauchy Data")
```

diffnormtol.int	<i>1-Sided Tolerance Limits for the Distribution of the Difference Between Two Independent Normal Random Variables</i>
-----------------	--

Description

Provides 1-sided tolerance limits for the difference between two independent normal random variables. If the ratio of the variances is known, then an exact calculation is performed. Otherwise, approximation methods are implemented.

Usage

```
diffnormtol.int(x1, x2, var.ratio = NULL, alpha = 0.05,
                P = 0.99, method = c("HALL", "GK", "RG"))
```

Arguments

x1	A vector of sample data which is distributed according to a normal distribution (sample 1).
x2	Another vector of sample data which is distributed according to a normal distribution (sample 2). It can be of a different sample size than the sample specified by x1.
var.ratio	A specified, known value of the variance ratio (i.e., the ratio of the variance for population 1 to the variance of population 2). If NULL, then the variance ratio is estimated according to one of the three methods specified in the method argument.
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of the population to be covered by the tolerance limits.
method	The method for estimating the variance ratio. This only needs to be specified in the case when var.ratio is not NULL. "HALL" is Hall's method, which takes a bias-corrected version of the ratio between the sample variance for sample 1 to the sample variance for sample 2. "GK" is the Guo-Krishnamoorthy method, which first calculates a bias-corrected version of the ratio between the sample variance for sample 2 to the sample variance for sample 1. The resulting limit is

then compared to the limit from Hall's method and the most conservative limit is chosen. "RG" is the Reiser-Guttman method, which is a biased version of the variance ratio that is calculated by taking the sample variance for sample 1 to the sample variance for sample 2. Typically, Hall's method or the Guo-Krishnamoorthy method are preferred to the Reiser-Guttman method.

Details

Satterthwaite's approximation for the degrees of freedom is used when the variance ratio is unknown.

Value

`diffnormtol.int` returns a data frame with items:

<code>alpha</code>	The specified significance level.
<code>P</code>	The proportion of the population covered by this tolerance interval.
<code>diff.bar</code>	The difference between the sample means.
<code>1-sided.lower</code>	The 1-sided lower tolerance bound.
<code>1-sided.upper</code>	The 1-sided upper tolerance bound.

Note

Unlike other tolerance interval functions, the output from `diffnormtol.int` cannot be passed to `plottol`.

References

- Guo, H. and Krishnamoorthy, K. (2004), New Approximate Inferential Methods for the Reliability Parameter in a Stress-Strength Model: The Normal Case, *Communications in Statistics - Theory and Methods*, **33**, 1715–1731.
- Hall, I. J. (1984), Approximate One-Sided Tolerance Limits for the Difference or Sum of Two Independent Normal Variates, *Journal of Quality Technology*, **16**, 15–19.
- Krishnamoorthy, K. and Mathew, T. (2009), *Statistical Tolerance Regions: Theory, Applications, and Computation*, Wiley.
- Reiser, B. J. and Guttman, I. (1986), Statistical Inference for $\Pr(Y < X)$: The Normal Case, *Technometrics*, **28**, 253–257.

See Also

[Normal, K. factor, normtol.int](#)

Examples

```
## 90%/99% tolerance limits for the difference between two
## simulated normal data sets. This data is taken from
## Krishnamoorthy and Mathew (2009). Note that there is a
```

```
## calculational error in their example, which yields different
## results with the output below.

x1 <- c(10.166, 5.889, 8.258, 7.303, 8.757)
x2 <- c(-0.204, 2.578, 1.182, 1.892, 0.786, -0.517, 1.156,
        0.980, 0.323, 0.437, 0.397, 0.050, 0.812, 0.720)

diffnormtol.int(x1, x2, alpha = 0.10, P = 0.99, method = "HALL")
diffnormtol.int(x1, x2, alpha = 0.10, P = 0.99, method = "GK")
diffnormtol.int(x1, x2, alpha = 0.10, P = 0.99, method = "RG")
diffnormtol.int(x1, x2, var.ratio = 3.8, alpha = 0.10, P = 0.99)
```

DiffProp

*Difference Between Two Proportions Distribution***Description**

Density (mass), distribution function, quantile function, and random generation for the difference between two proportions. This is determined by taking the difference between two independent beta distributions.

Usage

```
ddiffprop(x, k1, k2, n1, n2, a1 = 0.5, a2 = 0.5,
          log = FALSE, ...)
pdiffprop(q, k1, k2, n1, n2, a1 = 0.5, a2 = 0.5,
          lower.tail = TRUE, log.p = FALSE, ...)
qdiffprop(p, k1, k2, n1, n2, a1 = 0.5, a2 = 0.5,
          lower.tail = TRUE, log.p = FALSE, ...)
rdiffprop(n, k1, k2, n1, n2, a1 = 0.5, a2 = 0.5)
```

Arguments

x, q	Vector of quantiles.
p	Vector of probabilities.
n	The number of observations. If length>1, then the length is taken to be the number required.
k1, k2	The number of successes drawn from groups 1 and 2, respectively.
n1, n2	The sample sizes for groups 1 and 2, respectively.
a1, a2	The shift parameters for the beta distributions. For the fiducial approach, we know that the lower and upper limits are set at $a_1 = a_2 = 0$ and $a_1 = a_2 = 1$, respectively, for the true p_1 and p_2 . While computations can be performed on real values outside the unit interval, a warning message will be returned if such values are specified. For practical purposes, the default value of 0.5 should be used for each parameter.
log, log.p	Logical vectors. If TRUE, then the probabilities are given as $\log(p)$.
lower.tail	Logical vector. If TRUE, then probabilities are $P[X \leq x]$, else $P[X > x]$.
...	Additional arguments passed to the Appell F1 function.

Details

The difference between two proportions distribution has a fairly complicated functional form. Please see the article by Chen and Luo (2011), who corrected a typo in the article by Nadarajah and Kotz (2007), for the functional form of this distribution.

Value

ddiffprop gives the density (mass), pdiffprop gives the distribution function, qdiffprop gives the quantile function, and rdiffprop generates random deviates.

References

Chen, Y. and Luo, S. (2011), A Few Remarks on 'Statistical Distribution of the Difference of Two Proportions', *Statistics in Medicine*, **30**, 1913–1915.

Nadarajah, S. and Kotz, S. (2007), Statistical Distribution of the Difference of Two Proportions, *Statistics in Medicine*, **26**, 3518–3523.

See Also

[runif](#) and [.Random.seed](#) about random number generation.

Examples

```
## Randomly generated data from the difference between
## two proportions distribution.

set.seed(100)
x <- rdiffprop(n = 100, k1 = 2, k2 = 10, n1 = 17, n2 = 13)
hist(x, main = "Randomly Generated Data", prob = TRUE)

x.1 <- sort(x)
y <- ddiffprop(x = x.1, k1 = 2, k2 = 10, n1 = 17, n2 = 13)
lines(x.1, y, col = 2, lwd = 2)

plot(x.1, pdiffprop(q = x.1, k1 = 2, k2 = 10, n1 = 17,
  n2 = 13), type = "l", xlab = "x",
  ylab = "Cumulative Probabilities")

qdiffprop(p = 0.20, k1 = 2, k2 = 10, n1 = 17, n2 = 13,
  lower.tail = FALSE)
qdiffprop(p = 0.80, k1 = 2, k2 = 10, n1 = 17, n2 = 13)
```

Description

When providing two of the three quantities n , α , and P , this function solves for the third quantity in the context of distribution-free tolerance intervals.

Usage

```
distfree.est(n = NULL, alpha = NULL, P = NULL, side = 1)
```

Arguments

n	The necessary sample size to cover a proportion P of the population with confidence $1-\alpha$. Can be a vector.
α	1 minus the confidence level attained when it is desired to cover a proportion P of the population and a sample size n is provided. Can be a vector.
P	The proportion of the population to be covered with confidence $1-\alpha$ when a sample size n is provided. Can be a vector.
side	Whether a 1-sided or 2-sided tolerance interval is assumed (determined by $side = 1$ or $side = 2$, respectively).

Value

When providing two of the three quantities n , α , and P , `distfree.est` returns the third quantity. If more than one value of a certain quantity is specified, then a table will be returned.

References

Natrella, M. G. (1963), *Experimental Statistics: National Bureau of Standards - Handbook No. 91*, United States Government Printing Office, Washington, D.C.

See Also

[nptol.int](#)

Examples

```
## Solving for 1 minus the confidence level.  
distfree.est(n = 59, P = 0.95, side = 1)  
  
## Solving for the sample size.  
distfree.est(alpha = 0.05, P = 0.95, side = 1)  
  
## Solving for the proportion of the population to cover.  
distfree.est(n = 59, alpha = 0.05, side = 1)  
  
## Solving for sample sizes for many tolerance specifications.
```



```
distfree.est(alpha = seq(0.01, 0.05, 0.01),
             P = seq(0.80, 0.99, 0.01), side = 2)
```

 exp2tol.int

2-Parameter Exponential Tolerance Intervals

Description

Provides 1-sided or 2-sided tolerance intervals for data distributed according to a 2-parameter exponential distribution. Data with Type II censoring is permitted.

Usage

```
exp2tol.int(x, alpha = 0.05, P = 0.99, side = 1,
            method = c("GPU", "DUN", "KM"), type.2 = FALSE)
```

Arguments

x	A vector of data which is distributed according to the 2-parameter exponential distribution.
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of the population to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by side = 1 or side = 2, respectively).
method	The method for how the upper tolerance bound is approximated. "GPU" is the Guenther-Patil-Uppuluri method. "DUN" is the Dunsmore method, which has been empirically shown to be an improvement for samples greater than or equal to 8. "KM" is the Krishnamoorthy-Mathew method, which is typically more liberal than the other methods. More information on these methods can be found in the "References", which also highlight general sample size conditions as to when these different methods should be used.
type.2	Select TRUE if Type II censoring is present (i.e., the data set is censored at the maximum value present). The default is FALSE.

Value

exp2tol.int returns a data frame with items:

alpha	The specified significance level.
P	The proportion of the population covered by this tolerance interval.
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

References

- Dunsmore, I. R. (1978), Some Approximations for Tolerance Factors for the Two Parameter Exponential Distribution, *Technometrics*, **20**, 317–318.
- Engelhardt, M. and Bain, L. J. (1978), Tolerance Limits and Confidence Limits on Reliability for the Two-Parameter Exponential Distribution, *Technometrics*, **20**, 37–39.
- Guenther, W. C., Patil, S. A., and Uppuluri, V. R. R. (1976), One-Sided β -Content Tolerance Factors for the Two Parameter Exponential Distribution, *Technometrics*, **18**, 333–340.
- Krishnamoorthy, K. and Mathew, T. (2009), *Statistical Tolerance Regions: Theory, Applications, and Computation*, Wiley.

See Also

[TwoParExponential](#)

Examples

```
## 95%/90% 1-sided 2-parameter exponential tolerance intervals
## for a sample of size 50.

set.seed(100)
x <- r2exp(50, 6, shift = 55)
out <- exp2tol.int(x = x, alpha = 0.05, P = 0.90, side = 1,
                  method = "DUN", type.2 = FALSE)

out

plottol(out, x, plot.type = "both", side = "upper",
        x.lab = "2-Parameter Exponential Data")
```

exptol.int

Exponential Tolerance Intervals

Description

Provides 1-sided or 2-sided tolerance intervals for data distributed according to an exponential distribution. Data with Type II censoring is permitted.

Usage

```
exptol.int(x, alpha = 0.05, P = 0.99, side = 1, type.2 = FALSE)
```

Arguments

- | | |
|-------|---|
| x | A vector of data which is distributed according to an exponential distribution. |
| alpha | The level chosen such that 1-alpha is the confidence level. |
| P | The proportion of the population to be covered by this tolerance interval. |

side	Whether a 1-sided or 2-sided tolerance interval is required (determined by side = 1 or side = 2, respectively).
type.2	Select TRUE if Type II censoring is present (i.e., the data set is censored at the maximum value present). The default is FALSE.

Value

exptol.int returns a data frame with items:

alpha	The specified significance level.
P	The proportion of the population covered by this tolerance interval.
lambda.hat	The mean of the data (i.e., 1/rate).
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

References

Blischke, W. R. and Murthy, D. N. P. (2000), *Reliability: Modeling, Prediction, and Optimization*, John Wiley & Sons, Inc.

See Also

[Exponential](#)

Examples

```
## 95%/99% 1-sided exponential tolerance intervals for a
## sample of size 50.

set.seed(100)
x <- rexp(100, 0.004)
out <- exptol.int(x = x, alpha = 0.05, P = 0.99, side = 1,
                 type.2 = FALSE)
out

plottol(out, x, plot.type = "both", side = "lower",
        x.lab = "Exponential Data")
```

 exttol.int

Weibull (or Extreme-Value) Tolerance Intervals

Description

Provides 1-sided or 2-sided tolerance intervals for data distributed according to either a Weibull distribution or extreme-value (also called Gumbel) distributions.

Usage

```
exttol.int(x, alpha = 0.05, P = 0.99, side = 1,
           dist = c("Weibull", "Gumbel"), ext = c("min", "max"),
           NR.delta = 1e-8)
```

Arguments

x	A vector of data which is distributed according to either a Weibull distribution or an extreme-value distribution.
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of the population to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by side = 1 or side = 2, respectively).
dist	Select either dist = "Weibull" or dist = "Gumbel" if the data is distributed according to the Weibull or extreme-value distribution, respectively.
ext	If dist = "Gumbel", then select which extreme is to be modeled for the Gumbel distribution. The Gumbel distribution for the minimum (i.e., ext = "min") corresponds to a left-skewed distribution and the Gumbel distribution for the maximum (i.e., ext = "max") corresponds to a right-skewed distribution
NR.delta	The stopping criterion used for the Newton-Raphson algorithm when finding the maximum likelihood estimates of the Weibull or extreme-value distribution.

Details

Recall that the relationship between the Weibull distribution and the extreme-value distribution for the minimum is that if the random variable X is distributed according to a Weibull distribution, then the random variable $Y = \ln(X)$ is distributed according to an extreme-value distribution for the minimum.

If dist = "Weibull", then the natural logarithm of the data are taken so that a Newton-Raphson algorithm can be employed to find the MLEs of the extreme-value distribution for the minimum and then the data and MLEs are transformed back appropriately. No transformation is performed if dist = "Gumbel". The Newton-Raphson algorithm is initialized by the method of moments estimators for the parameters.

Value

exttol.int returns a data frame with items:

alpha	The specified significance level.
P	The proportion of the population covered by this tolerance interval.
shape.1	MLE for the shape parameter if dist = "Weibull" or for the location parameter if dist = "Gumbel".
shape.2	MLE for the scale parameter if dist = "Weibull" or dist = "Gumbel".
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

References

Bain, L. J. and Engelhardt, M. (1981), Simple Approximate Distributional Results for Confidence and Tolerance Limits for the Weibull Distribution Based on Maximum Likelihood Estimators, *Technometrics*, **23**, 15–20.

Coles, S. (2001), *An Introduction to Statistical Modeling of Extreme Values*, Springer.

See Also

[Weibull](#)

Examples

```
## 85%/90% 1-sided Weibull tolerance intervals for a sample
## of size 150.

set.seed(100)
x <- rweibull(150, 3, 75)
out <- exttol.int(x = x, alpha = 0.15, P = 0.90, side = 1,
                 dist = "Weibull")
out

plottol(out, x, plot.type = "both", side = "lower",
        x.lab = "Weibull Data")
```

F1

Appell's F1 Hypergeometric Function

Description

The Appell function of the first kind, which is a two variable extension of the hypergeometric distribution.

Usage

`F1(a, b, b.prime, c, x, y, ...)`

Arguments

`a, b, b.prime, c`

Appropriate parameters for this function.

`x, y`

The inputted values to evaluate this function such that each is less than 1 in absolute value.

`...`

Additional arguments passed to the integrate function.

Value

F1 returns the simple integral result for the Appell function of the first kind with the arguments specified above.

Note

This function is solved by using a simple integral representation for real numbers. While all four of the Appell functions can be extended to the complex plane, this is not an option for this code.

References

Bailey, W. N. (1935), *Generalised Hypergeometric Series*, Cambridge University Press.

See Also

[DiffProp](#), [integrate](#)

Examples

```
## Sample calculation.
```

```
F1(a = 3, b = 4, b.prime = 5, c = 13, x = 0.2, y = 0.4)
```

fidbintol.int	<i>Fiducial-Based Tolerance Intervals for the Function of Two Binomial Proportions</i>
---------------	--

Description

Provides 1-sided or 2-sided tolerance intervals for the function of two binomial proportions using fiducial quantities.

Usage

```
fidbintol.int(x1, x2, n1, n2, m1 = NULL, m2 = NULL, FUN,
             alpha = 0.05, P = 0.99, side = 1, K = 1000,
             B = 1000)
```

Arguments

x1	A value of observed "successes" from group 1.
x2	A value of observed "successes" from group 2.
n1	The total number of trials for group 1.
n2	The total number of trials for group 2.
m1	The total number of future trials for group 1. If NULL, then it is set to n1.
m2	The total number of future trials for group 2. If NULL, then it is set to n2.
FUN	Any reasonable (and meaningful) function taking exactly two arguments that we are interested in constructing a tolerance interval on.
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of the population to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by side = 1 or side = 2, respectively).
K	The number of fiducial quantities to be generated. The number of iterations should be at least as large as the default value of 1000. See Details for the definition of the fiducial quantity for a binomial proportion.
B	The number of iterations used for the Monte Carlo algorithm which determines the tolerance limits. The number of iterations should be at least as large as the default value of 1000.

Details

If X is observed from a $Bin(n, p)$ distribution, then the fiducial quantity for p is $Beta(X + 0.5, n - X + 0.5)$.

Value

fidbintol.int returns a list with two items. The first item (tol.limits) is a data frame with the following items:

alpha	The specified significance level.
P	The proportion of the population covered by this tolerance interval.
fn.est	A point estimate of the functional form of interest using the maximum likelihood estimates calculated with the inputted values of x1, x2, n1, and n2.
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

The second item (fn) simply returns the functional form specified by FUN.

References

Clopper, C. J. and Pearson, E. S. (1934), The Use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial, *Biometrika*, **26**, 404–413.

Krishnamoorthy, K. and Lee, M. (2010), Inference for Functions of Parameters in Discrete Distributions Based on Fiducial Approach: Binomial and Poisson Cases, *Journal of Statistical Planning and Inference*, **140**, 1182–1192.

Mathew, T. and Young, D. S. (2013), Fiducial-Based Tolerance Intervals for Some Discrete Distributions, *Computational Statistics and Data Analysis*, **61**, 38–49.

See Also

[fidnegbintol.int](#), [fidpoistol.int](#)

Examples

```
## 95%/99% 1-sided and 2-sided tolerance intervals for
## the difference between binomial proportions.

set.seed(100)

p1 <- 0.2
p2 <- 0.4
n1 <- n2 <- 200
x1 <- rbinom(1, n1, p1)
x2 <- rbinom(1, n2, p2)
fun.ti <- function(x, y) x - y

fidbintol.int(x1, x2, n1, n2, m1 = 500, m2 = 500, FUN = fun.ti,
             alpha = 0.05, P = 0.99, side = 1)
fidbintol.int(x1, x2, n1, n2, m1 = 500, m2 = 500, FUN = fun.ti,
             alpha = 0.05, P = 0.99, side = 2)
```

fidnegbintol.int	<i>Fiducial-Based Tolerance Intervals for the Function of Two Negative Binomial Proportions</i>
------------------	---

Description

Provides 1-sided or 2-sided tolerance intervals for the function of two negative binomial proportions using fiducial quantities.

Usage

```
fidnegbintol.int(x1, x2, n1, n2, m1 = NULL, m2 = NULL, FUN,
                 alpha = 0.05, P = 0.99, side = 1, K = 1000,
                 B = 1000)
```

Arguments

x1	A value of observed "failures" from group 1.
x2	A value of observed "failures" from group 2.
n1	The target number of successes for group 1.
n2	The target number of successes for group 2.
m1	The total number of future trials for group 1. If NULL, then it is set to n1.
m2	The total number of future trials for group 2. If NULL, then it is set to n2.
FUN	Any reasonable (and meaningful) function taking exactly two arguments that we are interested in constructing a tolerance interval on.
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of the population to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by side = 1 or side = 2, respectively).
K	The number of fiducial quantities to be generated. The number of iterations should be at least as large as the default value of 1000. See Details for the definition of the fiducial quantity for a negative binomial proportion.
B	The number of iterations used for the Monte Carlo algorithm which determines the tolerance limits. The number of iterations should be at least as large as the default value of 1000.

Details

If X is observed from a $NegBin(n, p)$ distribution, then the fiducial quantity for p is $Beta(n, X + 0.5)$.

Value

fidnegbintol.int returns a list with two items. The first item (tol.limits) is a data frame with the following items:

alpha	The specified significance level.
P	The proportion of the population covered by this tolerance interval.
fn.est	A point estimate of the functional form of interest using the maximum likelihood estimates calculated with the inputted values of x1, x2, n1, and n2.
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

The second item (fn) simply returns the functional form specified by FUN.

References

Cai, Y. and Krishnamoorthy, K. (2005), A Simple Improved Inferential Method for Some Discrete Distributions, *Computational Statistics and Data Analysis*, **48**, 605–621.

Clopper, C. J. and Pearson, E. S. (1934), The Use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial, *Biometrika*, **26**, 404–413.

Krishnamoorthy, K. and Lee, M. (2010), Inference for Functions of Parameters in Discrete Distributions Based on Fiducial Approach: Binomial and Poisson Cases, *Journal of Statistical Planning and Inference*, **140**, 1182–1192.

Mathew, T. and Young, D. S. (2013), Fiducial-Based Tolerance Intervals for Some Discrete Distributions, *Computational Statistics and Data Analysis*, **61**, 38–49.

See Also

[fidbintol.int](#), [fidpoistol.int](#)

Examples

```
## 95%/99% 1-sided and 2-sided tolerance intervals for
## the ratio of odds ratios for negative binomial proportions.

set.seed(100)

p1 <- 0.6
p2 <- 0.2
n1 <- n2 <- 50
x1 <- rnbinom(1, n1, p1)
x2 <- rnbinom(1, n2, p2)
fun.ti <- function(x, y) x * (1 - y) / (y * (1 - x))

fidnegbintol.int(x1, x2, n1, n2, m1 = 50, m2 = 50, FUN = fun.ti,
  alpha = 0.05, P = 0.99, side = 1)
fidnegbintol.int(x1, x2, n1, n2, m1 = 50, m2 = 50, FUN = fun.ti,
```

```
alpha = 0.05, P = 0.99, side = 2)
```

fidpoistol.int	<i>Fiducial-Based Tolerance Intervals for the Function of Two Poisson Rates</i>
----------------	---

Description

Provides 1-sided or 2-sided tolerance intervals for the function of two Poisson rates using fiducial quantities.

Usage

```
fidpoistol.int(x1, x2, n1, n2, m1 = NULL, m2 = NULL, FUN,
              alpha = 0.05, P = 0.99, side = 1, K = 1000,
              B = 1000)
```

Arguments

x1	A value of observed counts from group 1.
x2	A value of observed counts from group 2.
n1	The length of time that x1 was recorded over.
n2	The length of time that x2 was recorded over.
m1	The total number of future trials for group 1. If NULL, then it is set to n1.
m2	The total number of future trials for group 2. If NULL, then it is set to n2.
FUN	Any reasonable (and meaningful) function taking exactly two arguments that we are interested in constructing a tolerance interval on.
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of the population to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by side = 1 or side = 2, respectively).
K	The number of fiducial quantities to be generated. The number of iterations should be at least as large as the default value of 1000. See Details for the definition of the fiducial quantity for a Poisson rate.
B	The number of iterations used for the Monte Carlo algorithm which determines the tolerance limits. The number of iterations should be at least as large as the default value of 1000.

Details

If X is observed from a $Poi(n * \lambda)$ distribution, then the fiducial quantity for λ is $\chi^2_{2*x+1}/(2 * n)$.

Value

fidpoistol.int returns a list with two items. The first item (tol.limits) is a data frame with the following items:

alpha	The specified significance level.
P	The proportion of the population covered by this tolerance interval.
fn.est	A point estimate of the functional form of interest using the maximum likelihood estimates calculated with the inputted values of x1, x2, n1, and n2.
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

The second item (fn) simply returns the functional form specified by FUN.

References

- Cox, D. R. (1953), Some Simple Approximate Tests for Poisson Variates, *Biometrika*, **40**, 354–360.
- Krishnamoorthy, K. and Lee, M. (2010), Inference for Functions of Parameters in Discrete Distributions Based on Fiducial Approach: Binomial and Poisson Cases, *Journal of Statistical Planning and Inference*, **140**, 1182–1192.
- Mathew, T. and Young, D. S. (2013), Fiducial-Based Tolerance Intervals for Some Discrete Distributions, *Computational Statistics and Data Analysis*, **61**, 38–49.

See Also

[fidbintol.int](#), [fidnegbintol.int](#)

Examples

```
## 95%/99% 1-sided and 2-sided tolerance intervals for
## the ratio of two Poisson rates.

set.seed(100)

lambda1 <- 10
lambda2 <- 2
n1 <- 3000
n2 <- 3250
x1 <- rpois(1, n1 * lambda1)
x2 <- rpois(1, n2 * lambda2)
fun.ti <- function(x, y) x / y

fidpoistol.int(x1, x2, n1, n2, m1 = 2000, m2 = 2500,
              FUN = fun.ti, alpha = 0.05, P = 0.99, side = 1)
fidpoistol.int(x1, x2, n1, n2, m1 = 2000, m2 = 2500,
              FUN = fun.ti, alpha = 0.05, P = 0.99, side = 2)
```

gamtol.int

*Gamma (or Log-Gamma) Tolerance Intervals***Description**

Provides 1-sided or 2-sided tolerance intervals for data distributed according to either a gamma distribution or log-gamma distribution.

Usage

```
gamtol.int(x, alpha = 0.05, P = 0.99, side = 1,
          method = c("HE", "HE2", "WBE", "ELL", "KM", "EXACT",
                    "OCT"), m = 50, log.gamma = FALSE)
```

Arguments

x	A vector of data which is distributed according to either a gamma distribution or a log-gamma distribution.
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of the population to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by side = 1 or side = 2, respectively).
method	The method for calculating the k-factors. The k-factor for the 1-sided tolerance intervals is performed exactly and thus is the same for the chosen method. "HE" is the Howe method and is often viewed as being extremely accurate, even for small sample sizes. "HE2" is a second method due to Howe, which performs similarly to the Weissberg-Beatty method, but is computationally simpler. "WBE" is the Weissberg-Beatty method (also called the Wald-Wolfowitz method), which performs similarly to the first Howe method for larger sample sizes. "ELL" is the Ellison correction to the Weissberg-Beatty method when f is appreciably larger than n^2 . A warning message is displayed if f is not larger than n^2 . "KM" is the Krishnamoorthy-Mathew approximation to the exact solution, which works well for larger sample sizes. "EXACT" computes the k-factor exactly by finding the integral solution to the problem via the <code>integrate</code> function. Note the computation time of this method is largely determined by m. "OCT" is the Owen approach to compute the k-factor when controlling the tails so that there is not more than $(1-P)/2$ of the data in each tail of the distribution.
m	The maximum number of subintervals to be used in the <code>integrate</code> function. This is necessary only for method = "EXACT" and method = "OCT". The larger the number, the more accurate the solution. Too low of a value can result in an error. A large value can also cause the function to be slow for method = "EXACT".
log.gamma	If TRUE, then the data is considered to be from a log-gamma distribution, in which case the output gives tolerance intervals for the log-gamma distribution. The default is FALSE.

Details

Recall that if the random variable X is distributed according to a log-gamma distribution, then the random variable $Y = \ln(X)$ is distributed according to a gamma distribution.

Value

gamtol.int returns a data frame with items:

alpha	The specified significance level.
P	The proportion of the population covered by this tolerance interval.
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

References

Krishnamoorthy, K., Mathew, T., and Mukherjee, S. (2008), Normal-Based Methods for a Gamma Distribution: Prediction and Tolerance Intervals and Stress-Strength Reliability, *Technometrics*, **50**, 69–78.

See Also

[GammaDist](#), [K.factor](#)

Examples

```
## 99%/99% 1-sided gamma tolerance intervals for a sample
## of size 50.

set.seed(100)
x <- rgamma(50, 0.30, scale = 2)
out <- gamtol.int(x = x, alpha = 0.01, P = 0.99, side = 1,
                 method = "HE")
out

plottol(out, x, plot.type = "both", side = "upper",
        x.lab = "Gamma Data")
```

hypertol.int *Hypergeometric Tolerance Intervals*

Description

Provides 1-sided or 2-sided tolerance intervals for hypergeometric random variables. From a sampling without replacement perspective, these limits use the proportion of units from group A (e.g., "black balls" in an urn) in a sample to bound the number of potential units drawn from group A in a future sample taken from the universe.

Usage

```
hypertol.int(x, n, N, m = NULL, alpha = 0.05, P = 0.99,
            side = 1, method = c("EX", "LS", "CC"))
```

Arguments

x	The number of units from group A in the sample. Can be a vector, in which case the sum of x is used.
n	The size of the random sample of units selected.
N	The population size.
m	The quantity of units to be sampled from the universe for a future study. If m = NULL, then the tolerance limits will be constructed assuming n for this quantity.
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of units from group A in future samples of size m to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by side = 1 or side = 2, respectively).
method	The method for calculating the lower and upper confidence bounds, which are used in the calculation of the tolerance bounds. The default method is "EX", which is an exact-based method. "LS" is the large-sample method. "CC" gives a continuity-corrected version of the large-sample method.

Value

hypertol.int returns a data frame with items:

alpha	The specified significance level.
P	The proportion of units from group A in future samples of size m.
rate	The sampling rate determined by n/N.
p.hat	The proportion of units in the sample from group A, calculated by x/n.
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

Note

As this methodology is built using large-sample theory, if the sampling rate is less than 0.05, then a warning is generated stating that the results are not reliable. Also, compare the functionality of this procedure with the [acc.samp](#) procedure, which is to determine a minimal acceptance limit for a particular sampling plan.

References

- Brown, L. D., Cai, T. T., and DasGupta, A. (2001), Interval Estimation for a Binomial Proportion, *Statistical Science*, **16**, 101–133.
- Eichenberger, P., Hulliger, B., and Potterat, J. (2011), Two Measures for Sample Size Determination, *Survey Research Methods*, **5**, 27–37.
- Young, D. S. (2014), Tolerance Intervals for Hypergeometric and Negative Hypergeometric Variables, *Sankhya: The Indian Journal of Statistics, Series B*, to appear.

See Also

[acc.samp](#), [Hypergeometric](#)

Examples

```
## 90%/95% 1-sided and 2-sided hypergeometric tolerance
## intervals for a future sample of 30 when the universe
## is of size 100.

hypertol.int(x = 15, n = 50, N = 100, m = 30, alpha = 0.10,
             P = 0.95, side = 1, method = "LS")
hypertol.int(x = 15, n = 50, N = 100, m = 30, alpha = 0.10,
             P = 0.95, side = 1, method = "CC")
hypertol.int(x = 15, n = 50, N = 100, m = 30, alpha = 0.10,
             P = 0.95, side = 2, method = "LS")
hypertol.int(x = 15, n = 50, N = 100, m = 30, alpha = 0.10,
             P = 0.95, side = 2, method = "CC")
```

K.factor

Estimating K-factors for Tolerance Intervals Based on Normality

Description

Estimates k-factors for tolerance intervals based on normality.

Usage

```
K.factor(n, f = NULL, alpha = 0.05, P = 0.99, side = 1,
         method = c("HE", "HE2", "WBE", "ELL", "KM", "EXACT",
                    "OCT"), m = 50)
```


Arguments

n	The (effective) sample size.
f	The number of degrees of freedom associated with calculating the estimate of the population standard deviation. If NULL, then f is taken to be n-1.
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of the population to be covered by the tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by side = 1 or side = 2, respectively).
method	The method for calculating the k-factors. The k-factor for the 1-sided tolerance intervals is performed exactly and thus is the same for the chosen method. "HE" is the Howe method and is often viewed as being extremely accurate, even for small sample sizes. "HE2" is a second method due to Howe, which performs similarly to the Weissberg-Beatty method, but is computationally simpler. "WBE" is the Weissberg-Beatty method (also called the Wald-Wolfowitz method), which performs similarly to the first Howe method for larger sample sizes. "ELL" is the Ellison correction to the Weissberg-Beatty method when f is appreciably larger than n^2 . A warning message is displayed if f is not larger than n^2 . "KM" is the Krishnamoorthy-Mathew approximation to the exact solution, which works well for larger sample sizes. "EXACT" computes the k-factor exactly by finding the integral solution to the problem via the integrate function. Note the computation time of this method is largely determined by m. "OCT" is the Owen approach to compute the k-factor when controlling the tails so that there is not more than $(1-P)/2$ of the data in each tail of the distribution.
m	The maximum number of subintervals to be used in the integrate function. This is necessary only for method = "EXACT" and method = "OCT". The larger the number, the more accurate the solution. Too low of a value can result in an error. A large value can also cause the function to be slow for method = "EXACT".

Value

K.factor returns the k-factor for tolerance intervals based on normality with the arguments specified above.

Note

For larger sample sizes, there may be some accuracy issues with the 1-sided calculation since it depends on the noncentral t-distribution. The code is primarily intended to be used for moderate values of the noncentrality parameter. It will not be highly accurate, especially in the tails, for large values. See [TDist](#) for further details.

References

- Ellison, B. E. (1964), On Two-Sided Tolerance Intervals for a Normal Distribution, *Annals of Mathematical Statistics*, **35**, 762–772.
- Howe, W. G. (1969), Two-Sided Tolerance Limits for Normal Populations - Some Improvements, *Journal of the American Statistical Association*, **64**, 610–620.

Krishnamoorthy, K. and Mathew, T. (2009), *Statistical Tolerance Regions: Theory, Applications, and Computation*, Wiley.

Odeh, R. E. and Owen, D. B. (1980), *Tables for Normal Tolerance Limits, Sampling Plans, and Screening*, Marcel-Dekker.

Owen, D. B. (1964), Controls of Percentages in Both Tails of the Normal Distribution, *Technometrics*, **6**, 377-387.

Wald, A. and Wolfowitz, J. (1946), Tolerance Limits for a Normal Distribution, *Annals of the Mathematical Statistics*, **17**, 208–215.

Weissberg, A. and Beatty, G. (1969), Tables of Tolerance Limit Factors for Normal Distributions, *Technometrics*, **2**, 483–500.

See Also

[integrate](#), [K.table](#), [normtol.int](#), [TDist](#)

Examples

```
## Showing the effect of the Howe, Weissberg-Beatty,
## and exact estimation methods as the sample size increases.

K.factor(10, P = 0.95, side = 2, method = "HE")
K.factor(10, P = 0.95, side = 2, method = "WBE")
K.factor(10, P = 0.95, side = 2, method = "EXACT", m = 50)

K.factor(100, P = 0.95, side = 2, method = "HE")
K.factor(100, P = 0.95, side = 2, method = "WBE")
K.factor(100, P = 0.95, side = 2, method = "EXACT", m = 50)
```

K.table

Tables of K-factors for Tolerance Intervals Based on Normality

Description

Tabulated summary of k-factors for tolerance intervals based on normality. The user can specify multiple values for each of the three inputs.

Usage

```
K.table(n, alpha, P, side = 1, f = NULL,
        by.arg = c("n", "alpha", "P"))
```

Arguments

n	A vector of (effective) sample sizes.
alpha	The level chosen such that 1-alpha is the confidence level. Can be a vector.
P	The proportion of the population to be covered by this tolerance interval. Can be a vector.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by side = 1 or side = 2, respectively).
f	The number of degrees of freedom associated with calculating the estimate of the population standard deviation. If NULL, then f is taken to be n-1. Only a single value can be specified for f.
by.arg	How you would like the output organized. If by.arg = "n", then the output provides a list of matrices sorted by the values specified in n. The matrices have rows corresponding to the values specified by 1-alpha and columns corresponding to the values specified by P. If by.arg = "alpha", then the output provides a list of matrices sorted by the values specified in 1-alpha. The matrices have rows corresponding to the values specified by n and columns corresponding to the values specified by P. If by.arg = "P", then the output provides a list of matrices sorted by the values specified in P. The matrices have rows corresponding to the values specified by 1-alpha and columns corresponding to the values specified by n.

Details

The method used for estimating the k-factors is that due to Howe as it is generally viewed as more accurate than the Weissberg-Beatty method.

Value

K.table returns a list with a structure determined by the argument by.arg described above.

References

Howe, W. G. (1969), Two-Sided Tolerance Limits for Normal Populations - Some Improvements, *Journal of the American Statistical Association*, **64**, 610-620.

Weissberg, A. and Beatty, G. (1969), Tables of Tolerance Limit Factors for Normal Distributions, *Technometrics*, **2**, 483-500.

See Also

[K.factor](#)

Examples

```
## Tables generated for each value of the sample size.
K.table(n = seq(50, 100, 10), alpha = c(0.01, 0.05, 0.10),
        P = c(0.90, 0.95, 0.99), by.arg = "n")
```

```
## Tables generated for each value of the confidence level.

K.table(n = seq(50, 100, 10), alpha = c(0.01, 0.05, 0.10),
        P = c(0.90, 0.95, 0.99), by.arg = "alpha")

## Tables generated for each value of the coverage proportion.

K.table(n = seq(50, 100, 10), alpha = c(0.01, 0.05, 0.10),
        P = c(0.90, 0.95, 0.99), by.arg = "P")
```

laptol.int

Laplace Tolerance Intervals

Description

Provides 1-sided or 2-sided tolerance intervals for data distributed according to a Laplace distribution.

Usage

```
laptol.int(x, alpha = 0.05, P = 0.99, side = 1)
```

Arguments

x	A vector of data which is distributed according to a Laplace distribution.
alpha	The level chosen such that $1-\alpha$ is the confidence level.
P	The proportion of the population to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by $side = 1$ or $side = 2$, respectively).

Value

laptol.int returns a data frame with items:

alpha	The specified significance level.
P	The proportion of the population covered by this tolerance interval.
1-sided.lower	The 1-sided lower tolerance bound. This is given only if $side = 1$.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if $side = 1$.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if $side = 2$.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if $side = 2$.

References

Bain, L. J. and Engelhardt, M. (1973), Interval Estimation for the Two Parameter Double Exponential Distribution, *Technometrics*, **15**, 875–887.

Examples

```
## First generate data from a Laplace distribution with location
## parameter 70 and scale parameter 3.

set.seed(100)
tmp <- runif(40)
x <- rep(70, 40) - sign(tmp - 0.5)*rep(3, 40)*
      log(2*ifelse(tmp < 0.5, tmp, 1-tmp))

## 95%/90% 1-sided Laplace tolerance intervals for the sample
## of size 40 generated above.

out <- laptol.int(x = x, alpha = 0.05, P = 0.90, side = 1)
out

plottol(out, x, plot.type = "hist", side = "two",
        x.lab = "Laplace Data")
```

logistol.int

Logistic (or Log-Logistic) Tolerance Intervals

Description

Provides 1-sided or 2-sided tolerance intervals for data distributed according to a logistic or log-logistic distribution.

Usage

```
logistol.int(x, alpha = 0.05, P = 0.99, log.log = FALSE,
            side = 1)
```

Arguments

x	A vector of data which is distributed according to a logistic or log-logistic distribution.
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of the population to be covered by this tolerance interval.
log.log	If TRUE, then the data is considered to be from a log-logistic distribution, in which case the output gives tolerance intervals for the log-logistic distribution. The default is FALSE.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by side = 1 or side = 2, respectively).

Details

Recall that if the random variable X is distributed according to a log-logistic distribution, then the random variable $Y = \ln(X)$ is distributed according to a logistic distribution.

Value

logistol.int returns a data frame with items:

alpha	The specified significance level.
P	The proportion of the population covered by this tolerance interval.
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

References

- Balakrishnan, N. (1992), *Handbook of the Logistic Distribution*, Marcel Dekker, Inc.
- Hall, I. J. (1975), One-Sided Tolerance Limits for a Logistic Distribution Based on Censored Samples, *Biometrics*, **31**, 873–880.

See Also

[Logistic](#)

Examples

```
## 90%/95% 1-sided logistic tolerance intervals for a sample
## of size 20.

set.seed(100)
x <- rlogis(20, 5, 1)
out <- logistol.int(x = x, alpha = 0.10, P = 0.95,
                   log.log = FALSE, side = 1)

out

plottol(out, x, plot.type = "control", side = "two",
        x.lab = "Logistic Data")
```

mvregtol.region

Multivariate (Multiple) Linear Regression Tolerance Regions

Description

Determines the appropriate tolerance factor for computing multivariate (multiple) linear regression tolerance regions based on Monte Carlo simulation.

Usage

```
mvregtol.region(y, x, new.x = NULL, int = TRUE, alpha = 0.05,
               P = 0.99, B = 1000)
```

Arguments

y	An $n \times q$ matrix of responses assumed to be drawn from a q -dimensional multivariate normal distribution. n pertains to the sample size.
x	An $n \times m$ matrix of predictors (i.e., the design matrix). m pertains to the number of predictors. Do NOT include a column of 1's if assuming an intercept. This is controlled by the <code>int</code> argument.
new.x	A matrix of new values for which to approximate k -factors. This must be a matrix with m columns.
int	If <code>int = TRUE</code> , then an intercept is assumed.
alpha	The level chosen such that $1 - \alpha$ is the confidence level.
P	The proportion of the population to be covered by this tolerance region.
B	The number of iterations used for the Monte Carlo algorithm which determines the tolerance factor. The number of iterations should be at least as large as the default value of 1000.

Details

A basic sketch of how the algorithm works is as follows:

- (1) Generate independent chi-square random variables and Wishart random matrices.
- (2) Compute the eigenvalues of the randomly generated Wishart matrices.
- (3) Iterate the above steps to generate a set of B sample values such that the $100(1 - \alpha)$ -th percentile is an approximate tolerance factor.

Value

`mvregtol.region` returns a matrix where the first column is the k -factor, the next q columns are the estimated responses from the least squares fit, and the final m columns are the predictor values. The first n rows of the matrix pertain to the raw data as specified by `y` and `x`. If values for `new.x` are specified, then there is one additional row appended to this output for each row in the matrix `new.x`.

References

- Anderson, T. W. (2003) *An Introduction to Multivariate Statistical Analysis*, Third Edition, Wiley.
- Krishnamoorthy, K. and Mathew, T. (2009), *Statistical Tolerance Regions: Theory, Applications, and Computation*, Wiley.
- Krishnamoorthy, K. and Mondal, S. (2008), Tolerance Factors in Multiple and Multivariate Linear Regressions, *Communications in Statistics - Simulation and Computation*, **37**, 546–559.

Examples

```
## 95%/95% multivariate regression tolerance factors using
## a fertilizer data set presented in Anderson (2003, p. 374).

grain <- c(40, 17, 9, 15, 6, 12, 5, 9)
straw <- c(53, 19, 10, 29, 13, 27, 19, 30)
```

```
fert <- c(24, 11, 5, 12, 7, 14, 11, 18)
Y <- cbind(grain, straw)
X <- cbind(fert)
new.x <- c(10, 15, 20)

set.seed(100)
out <- mvregtol.region(Y, X, new.x, int = TRUE, alpha = 0.05,
                      P = 0.95, B = 5000)

out
```

mvtol.region

Multivariate Normal Tolerance Regions

Description

Determines the appropriate tolerance factor for computing multivariate normal tolerance regions based on Monte Carlo methods or other approximations.

Usage

```
mvtol.region(x, alpha = 0.05, P = 0.99, B = 1000, M = 1000,
             method = c("KM", "AM", "GM", "HM", "MHM", "V11",
                       "HM.V11", "MC"))
```

Arguments

x	An $n \times p$ matrix of data assumed to be drawn from a p -dimensional multivariate normal distribution. n pertains to the sample size.
alpha	The level chosen such that $1 - \alpha$ is the confidence level. A vector of alpha values may be specified.
P	The proportion of the population to be covered by this tolerance region. A vector of P values may be specified.
B	The number of iterations used for the Monte Carlo algorithms (i.e., when <code>method = "KM"</code> or <code>"MC"</code>), which determines the tolerance factor. The number of iterations should be at least as large as the default value of 1000.
M	The number of iterations used for the inner loop of the Monte Carlo algorithm specified through <code>method = "MC"</code> . The number of iterations should be at least as large as the default value of 1000. Note that this is not required for <code>method = "KM"</code> since that algorithm handles the eigenvalues differently in the estimation of the tolerance factor.
method	The method for estimating the tolerance factors. <code>"KM"</code> is the Krishnamoorthy-Mondal method, which is the method implemented in previous versions of the tolerance package. It is one of the more accurate methods available. <code>"AM"</code> is an approximation method based on the arithmetic mean. <code>"GM"</code> is an approximation method based on the geometric mean. <code>"HM"</code> is an approximation method based on the harmonic mean. <code>"MHM"</code> is a modified approach based on the harmonic mean. <code>"V11"</code> is a method that utilizes a certain partitioning of a Wishart

random matrix for deriving an approximate tolerance factor. "HM.V11" is a hybrid method of the "HM" and "V11" methods. "MC" is a simple Monte Carlo approach to estimating the tolerance factor, which is computationally expensive as the values of B and M increase.

Details

All of the methods are outlined in the references that we provided. In practice, we recommend using the Krishnamoorthy-Mondal approach. A basic sketch of how the Krishnamoorthy-Mondal algorithm works is as follows:

- (1) Generate independent chi-square random variables and Wishart random matrices.
- (2) Compute the eigenvalues of the randomly generated Wishart matrices.
- (3) Iterate the above steps to generate a set of B sample values such that the $100(1-\alpha)$ -th percentile is an approximate tolerance factor.

Value

mvtol.region returns a matrix where the rows pertain to each confidence level $1-\alpha$ specified and the columns pertain to each proportion level P specified.

References

Krishnamoorthy, K. and Mathew, T. (1999), Comparison of Approximation Methods for Computing Tolerance Factors for a Multivariate Normal Population, *Technometrics*, **41**, 234–249.

Krishnamoorthy, K. and Mondal, S. (2006), Improved Tolerance Factors for Multivariate Normal Distributions, *Communications in Statistics - Simulation and Computation*, **35**, 461–478.

Examples

```
## 90%/90% bivariate normal tolerance region.

set.seed(100)
x1 <- rnorm(100, 0, 0.2)
x2 <- rnorm(100, 0, 0.5)
x <- cbind(x1, x2)

out1 <- mvtol.region(x = x, alpha = 0.10, P = 0.90, B = 1000,
                    method = "KM")

out1
plottol(out1, x)

## 90%/90% trivariate normal tolerance region.

set.seed(100)
x1 <- rnorm(100, 0, 0.2)
x2 <- rnorm(100, 0, 0.5)
x3 <- rnorm(100, 5, 1)
x <- cbind(x1, x2, x3)
mvtol.region(x = x, alpha = c(0.10, 0.05, 0.01),
```

```

P = c(0.90, 0.95, 0.99), B = 1000, method = "KM")

out2 <- mvtol.region(x = x, alpha = 0.10, P = 0.90, B = 1000,
                    method = "KM")

out2
plottol(out2, x)

```

negbintol.int

Negative Binomial Tolerance Intervals

Description

Provides 1-sided or 2-sided tolerance intervals for negative binomial random variables. From a statistical quality control perspective, these limits use the number of failures that occur to reach n successes to bound the number of failures for a specified amount of future successes (m).

Usage

```

negbintol.int(x, n, m = NULL, alpha = 0.05, P = 0.99,
             side = 1, method = c("LS", "WU", "CB",
                                 "CS", "SC", "LR", "SP", "CC"))

```

Arguments

x	The total number of failures that occur from a sample of size n . Can be a vector of length n , in which case the sum of x is computed.
n	The target number of successes (sometimes called size) for each trial.
m	The target number of successes in a future lot for which the tolerance limits will be calculated. If $m = \text{NULL}$, then the tolerance limits will be constructed assuming n for the target number of future successes.
alpha	The level chosen such that $1 - \alpha$ is the confidence level.
P	The proportion of the defective (or acceptable) units in future samples of size m to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by $\text{side} = 1$ or $\text{side} = 2$, respectively).
method	The method for calculating the lower and upper confidence bounds, which are used in the calculation of the tolerance bounds. The default method is "LS", which is the large-sample method based on the MLE. "WU" is a Wald-type interval based on the UMVUE of the negative binomial proportion. "CB" is the Casella-Berger exact method. "CS" is a method based on chi-square percentiles. "SC" is the score method. "LR" is a likelihood ratio-based method. "SP" is a method using a saddlepoint approximation for the confidence intervals. "CC" gives a continuity-corrected version of the large-sample method and is appropriate when n is large. More information on these methods can be found in the "References".

Details

This function takes the approach for Poisson and binomial random variables developed in Hahn and Chandra (1981) and applies it to the negative binomial case.

Value

negbintol.int returns a data frame with items:

alpha	The specified significance level.
P	The proportion of defective (or acceptable) units in future samples of size m.
pi.hat	The probability of success in each trial, calculated by $n/(n+x)$.
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

Note

Recall that the geometric distribution is the negative binomial distribution where the size is 1. Therefore, the case when $n = m = 1$ will provide tolerance limits for a geometric distribution.

References

- Casella, G. and Berger, R. L. (1990), *Statistical Inference*, Duxbury Press.
- Hahn, G. J. and Chandra, R. (1981), Tolerance Intervals for Poisson and Binomial Variables, *Journal of Quality Technology*, **13**, 100–110.
- Tian, M., Tang, M. L., Ng, H. K. T., and Chan, P. S. (2009), A Comparative Study of Confidence Intervals for Negative Binomial Proportions, *Journal of Statistical Computation and Simulation*, **79**, 241–249.
- Young, D. S. (2014), A Procedure for Approximate Negative Binomial Tolerance Intervals, *Journal of Statistical Computation and Simulation*, **84**, 438–450.

See Also

[NegBinomial](#), [umatol.int](#)

Examples

```
## Comparison of 95%/99% 1-sided tolerance limits with
## 50 failures before 10 successes are reached.

negbintol.int(x = 50, n = 10, side = 1, method = "LS")
negbintol.int(x = 50, n = 10, side = 1, method = "WU")
negbintol.int(x = 50, n = 10, side = 1, method = "CB")
negbintol.int(x = 50, n = 10, side = 1, method = "CS")
negbintol.int(x = 50, n = 10, side = 1, method = "SC")
```

```

negbintol.int(x = 50, n = 10, side = 1, method = "LR")
negbintol.int(x = 50, n = 10, side = 1, method = "SP")
negbintol.int(x = 50, n = 10, side = 1, method = "CC")

## 95%/99% 1-sided tolerance limits and 2-sided tolerance
## interval for the same setting above, but when we are
## interested in a future experiment that requires 20 successes
## be reached for each trial.

negbintol.int(x = 50, n = 10, m = 20, side = 1)
negbintol.int(x = 50, n = 10, m = 20, side = 2)

```

NegHypergeometric *The Negative Hypergeometric Distribution*

Description

Density, distribution function, quantile function, and random generation for the negative hypergeometric distribution.

Usage

```

dnhyper(x, m, n, k, log = FALSE)
pnhyper(q, m, n, k, lower.tail = TRUE, log.p = FALSE)
qnhyper(p, m, n, k, lower.tail = TRUE, log.p = FALSE)
rnhyper(nn, m, n, k)

```

Arguments

<code>x, q</code>	Vector of quantiles representing the number of trials until k successes have occurred (e.g., until k white balls have been drawn from an urn without replacement).
<code>m</code>	The number of successes in the population (e.g., the number of white balls in the urn).
<code>n</code>	The population size (e.g., the total number of balls in the urn).
<code>k</code>	The number of successes (e.g., white balls) to achieve with the sample.
<code>p</code>	Vector of probabilities, which must be between 0 and 1.
<code>nn</code>	The number of observations. If <code>length>1</code> , then the length is taken to be the number required.
<code>log, log.p</code>	Logical vectors. If TRUE, then probabilities are given as $\log(p)$.
<code>lower.tail</code>	Logical vector. If TRUE, then probabilities are $P[X \leq x]$, else $P[X > x]$.

Details

A negative hypergeometric distribution (sometimes called the inverse hypergeometric distribution) models the total number of trials until k successes occur. Compare this to the negative binomial distribution, which models the number of failures that occur until a specified number of successes has been reached. The negative hypergeometric distribution has density

$$p(x) = \frac{\binom{x-1}{k-1} \binom{n-x}{m-k}}{\binom{n}{m}}$$

for $x = k, k + 1, \dots, n - m + k$.

Value

`dnhyper` gives the density, `pnhyper` gives the distribution function, `qnhyper` gives the quantile function, and `rnhyper` generates random deviates.

Invalid arguments will return value NaN, with a warning.

References

Wilks, S. S. (1963), *Mathematical Statistics*, Wiley.

See Also

[runif](#) and [.Random.seed](#) about random number generation.

Examples

```
## Randomly generated data from the negative hypergeometric
## distribution.

set.seed(100)
x <- rnhyper(nn = 1000, m = 15, n = 40, k = 10)
hist(x, main = "Randomly Generated Data", prob = TRUE)

x.1 = sort(x)
y <- dnhyper(x = x.1, m = 15, n = 40, k = 10)
lines(x.1, y, col = 2, lwd = 2)

plot(x.1, pnhyper(q = x.1, m = 15, n = 40, k = 10),
      type = "l", xlab = "x", ylab = "Cumulative Probabilities")

qnhyper(p = 0.20, m = 15, n = 40, k = 10, lower.tail = FALSE)
qnhyper(p = 0.80, m = 15, n = 40, k = 10)
```

neghypertol.int

*Negative Hypergeometric Tolerance Intervals***Description**

Provides 1-sided or 2-sided tolerance intervals for negative hypergeometric random variables. When sampling without replacement, these limits are on the total number of expected draws in a future sample in order to achieve a certain number from group A (e.g., "black balls" in an urn).

Usage

```
neghypertol.int(x, n, N, m = NULL, alpha = 0.05, P = 0.99,
               side = 1, method = c("EX", "LS", "CC"))
```

Arguments

x	The number of units drawn in order to achieve n successes. Can be a vector, in which case the sum of x is used.
n	The target number of successes in the sample drawn (e.g., the number of "black balls" you are to draw in the sample).
N	The population size (e.g., the total number of balls in the urn).
m	The target number of successes to be sampled from the universe for a future study. If m = NULL, then the tolerance limits will be constructed assuming n for this quantity.
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of units from group A in future samples of size m to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by side = 1 or side = 2, respectively).
method	The method for calculating the lower and upper confidence bounds, which are used in the calculation of the tolerance bounds. The default method is "EX", which is an exact-based method. "LS" is the large-sample method. "CC" gives a continuity-corrected version of the large-sample method.

Value

neghypertol.int returns a data frame with items:

alpha	The specified significance level.
P	The proportion of units from group A in future samples of size m.
rate	The sampling rate determined by x/N.
p.hat	The proportion of units in the sample from group A, calculated by n/x.
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

Note

As this methodology is built using large-sample theory, if the sampling rate is less than 0.05, then a warning is generated stating that the results are not reliable.

References

Khan, R. A. (1994), A Note on the Generating Function of a Negative Hypergeometric Distribution, *Sankhya: The Indian Journal of Statistics, Series B*, **56**, 309–313.

Young, D. S. (2014), Tolerance Intervals for Hypergeometric and Negative Hypergeometric Variables, *Sankhya: The Indian Journal of Statistics, Series B*, to appear.

See Also

[acc.samp](#), [NegHypergeometric](#)

Examples

```
## 90%/95% 1-sided and 2-sided negative hypergeometric
## tolerance intervals for a future number of 300 successes
## when the universe is of size 1000. The estimates are
## based on having drawn 425 in another sample to achieve
## 200 successes.
```

```
neghypertol.int(425, 200, 1000, m = 300, alpha = 0.05,
                P = 0.95, side = 1, method = "LS")
neghypertol.int(425, 200, 1000, m = 300, alpha = 0.05,
                P = 0.95, side = 1, method = "CC")
```

```
neghypertol.int(425, 200, 1000, m = 300, alpha = 0.05,
                P = 0.95, side = 2, method = "LS")
neghypertol.int(425, 200, 1000, m = 300, alpha = 0.05,
                P = 0.95, side = 2, method = "CC")
```

nlregtol.int

Nonlinear Regression Tolerance Bounds

Description

Provides 1-sided or 2-sided nonlinear regression tolerance bounds.

Usage

```
nlregtol.int(formula, xy.data = data.frame(), x.new = NULL,
             side = 1, alpha = 0.05, P = 0.99, maxiter = 50,
             ...)
```

Arguments

formula	A nonlinear model formula including variables and parameters.
xy.data	A data frame in which to evaluate the formulas in formula. The first column of xy.data must be the response variable.
x.new	Any new levels of the predictor(s) for which to report the tolerance bounds. The number of columns must be 1 less than the number of columns for xy.data.
side	Whether a 1-sided or 2-sided tolerance bound is required (determined by side = 1 or side = 2, respectively).
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of the population to be covered by the tolerance bound(s).
maxiter	A positive integer specifying the maximum number of iterations that the nonlinear least squares routine (nls) should run.
...	Optional arguments passed to nls when estimating the nonlinear regression equation.

Details

It is highly recommended that the user specify starting values for the nls routine.

Value

nlregtol.int returns a data frame with items:

alpha	The specified significance level.
P	The proportion of the population covered by the tolerance bound(s).
y.hat	The predicted value of the response for the fitted nonlinear regression model.
y	The value of the response given in the first column of xy.data. This data frame is sorted by this value.
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

References

- Wallis, W. A. (1951), Tolerance Intervals for Linear Regression, in *Second Berkeley Symposium on Mathematical Statistics and Probability*, ed. J. Neyman, Berkeley: University of CA Press, 43–51.
- Young, D. S. (2013), Regression Tolerance Intervals, *Communications in Statistics - Simulation and Computation*, **42**, 2040–2055.

See Also

[nls](#)

Examples

```
## 95%/95% 2-sided nonlinear regression tolerance bounds
## for a sample of size 50.

set.seed(100)
x <- runif(50, 5, 45)
f1 <- function(x, b1, b2) b1 + (0.49 - b1)*exp(-b2*(x - 8)) +
  rnorm(50, sd = 0.01)
y <- f1(x, 0.39, 0.11)
formula <- as.formula(y ~ b1 + (0.49 - b1)*exp(-b2*(x - 8)))
out <- nlregtol.int(formula = formula,
  xy.data = data.frame(cbind(y, x)),
  x.new=cbind(c(10, 20)), side = 2,
  alpha = 0.05, P = 0.95)

out

plottol(out, x = x, y = y, side = "two", x.lab = "X",
  y.lab = "Y")
```

normtol.int

*Normal (or Log-Normal) Tolerance Intervals***Description**

Provides 1-sided or 2-sided tolerance intervals for data distributed according to either a normal distribution or log-normal distribution.

Usage

```
normtol.int(x, alpha = 0.05, P = 0.99, side = 1,
  method = c("HE", "HE2", "WBE", "ELL", "KM",
  "EXACT", "OCT"), m = 50, log.norm = FALSE)
```

Arguments

x	A vector of data which is distributed according to either a normal distribution or a log-normal distribution.
alpha	The level chosen such that $1-\alpha$ is the confidence level.
P	The proportion of the population to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by $side = 1$ or $side = 2$, respectively).
method	The method for calculating the k-factors. The k-factor for the 1-sided tolerance intervals is performed exactly and thus is the same for the chosen method. "HE" is the Howe method and is often viewed as being extremely accurate, even

for small sample sizes. "HE2" is a second method due to Howe, which performs similarly to the Weissberg-Beatty method, but is computationally simpler. "WBE" is the Weissberg-Beatty method (also called the Wald-Wolfowitz method), which performs similarly to the first Howe method for larger sample sizes. "ELL" is the Ellison correction to the Weissberg-Beatty method when f is appreciably larger than n^2 . A warning message is displayed if f is not larger than n^2 . "KM" is the Krishnamoorthy-Mathew approximation to the exact solution, which works well for larger sample sizes. "EXACT" computes the k-factor exactly by finding the integral solution to the problem via the `integrate` function. Note the computation time of this method is largely determined by m . "OCT" is the Owen approach to compute the k-factor when controlling the tails so that there is not more than $(1-P)/2$ of the data in each tail of the distribution.

<code>m</code>	The maximum number of subintervals to be used in the <code>integrate</code> function. This is necessary only for <code>method = "EXACT"</code> and <code>method = "OCT"</code> . The larger the number, the more accurate the solution. Too low of a value can result in an error. A large value can also cause the function to be slow for <code>method = "EXACT"</code> .
<code>log.norm</code>	If TRUE, then the data is considered to be from a log-normal distribution, in which case the output gives tolerance intervals for the log-normal distribution. The default is FALSE.

Details

Recall that if the random variable X is distributed according to a log-normal distribution, then the random variable $Y = \ln(X)$ is distributed according to a normal distribution.

Value

`normtol.int` returns a data frame with items:

<code>alpha</code>	The specified significance level.
<code>P</code>	The proportion of the population covered by this tolerance interval.
<code>x.bar</code>	The sample mean.
<code>1-sided.lower</code>	The 1-sided lower tolerance bound. This is given only if <code>side = 1</code> .
<code>1-sided.upper</code>	The 1-sided upper tolerance bound. This is given only if <code>side = 1</code> .
<code>2-sided.lower</code>	The 2-sided lower tolerance bound. This is given only if <code>side = 2</code> .
<code>2-sided.upper</code>	The 2-sided upper tolerance bound. This is given only if <code>side = 2</code> .

References

- Howe, W. G. (1969), Two-Sided Tolerance Limits for Normal Populations - Some Improvements, *Journal of the American Statistical Association*, **64**, 610–620.
- Wald, A. and Wolfowitz, J. (1946), Tolerance Limits for a Normal Distribution, *Annals of Mathematical Statistics*, **17**, 208–215.
- Weissberg, A. and Beatty, G. (1969), Tables of Tolerance Limit Factors for Normal Distributions, *Technometrics*, **2**, 483–500.

See Also[Normal, K. factor](#)**Examples**

```
## 95%/95% 2-sided normal tolerance intervals for a sample
## of size 100.

set.seed(100)
x <- rnorm(100, 0, 0.2)
out <- normtol.int(x = x, alpha = 0.05, P = 0.95, side = 2,
                  method = "HE", log.norm = FALSE)

out

plottol(out, x, plot.type = "both", side = "two",
        x.lab = "Normal Data")
```

np.order

Sample Size Determination for Tolerance Limits Based on Order Statistics

Description

For given values of m , α , and P , this function solves the necessary sample size such that the r -th (or $(n-s+1)$ -th) order statistic is the $[100(1-\alpha)\%, 100(P)\%]$ lower (or upper) tolerance limit (see the Details section below for further explanation). This function can also report all combinations of order statistics for 2-sided intervals.

Usage

```
np.order(m, alpha = 0.05, P = 0.99, indices = FALSE)
```

Arguments

<code>m</code>	See the Details section below for how m is defined.
<code>alpha</code>	1 minus the confidence level attained when it is desired to cover a proportion P of the population with the order statistics.
<code>P</code>	The proportion of the population to be covered with confidence $1-\alpha$ with the order statistics.
<code>indices</code>	An optional argument to report all combinations of order statistics indices for the upper and lower limits of the 2-sided intervals. Note that this can only be calculated when $m>1$.

Details

For the 1-sided tolerance limits, $m=s+r$ such that the probability is at least $1-\alpha$ that at least the proportion P of the population is below the $(n-s+1)$ -th order statistic for the upper limit or above the r -th order statistic for the lower limit. This means for the 1-sided upper limit that $r=1$, while for the 1-sided lower limit it means that $s=1$. For the 2-sided tolerance intervals, $m=s+r$ such that the probability is at least $1-\alpha$ that at least the proportion P of the population is between the r -th and $(n-s+1)$ -th order statistics. Thus, all combinations of $r>0$ and $s>0$ such that $m=s+r$ are considered.

Value

If `indices = FALSE`, then a single number is returned for the necessary sample size such that the r -th (or $(n-s+1)$ -th) order statistic is the $[100(1-\alpha)\%, 100(P)\%]$ lower (or upper) tolerance limit. If `indices = TRUE`, then a list is returned with a single number for the necessary sample size and a matrix with 2 columns where each row gives the pairs of indices for the order statistics for all permissible $[100(1-\alpha)\%, 100(P)\%]$ 2-sided tolerance intervals.

References

Hanson, D. L. and Owen, D. B. (1963), Distribution-Free Tolerance Limits Elimination of the Requirement That Cumulative Distribution Functions Be Continuous, *Technometrics*, **5**, 518–522.

Scheffe, H. and Tukey, J. W. (1945), Non-Parametric Estimation I. Validation of Order Statistics, *Annals of Mathematical Statistics*, **16**, 187–192.

See Also

[nptol.int](#)

Examples

```
## Only requesting the sample size.
np.order(m = 5, alpha = 0.05, P = 0.95)

## Requesting the order statistics indices as well.
np.order(m = 5, alpha = 0.05, P = 0.95, indices = TRUE)
```

npregtol.int

Nonparametric Regression Tolerance Bounds

Description

Provides 1-sided or 2-sided nonparametric regression tolerance bounds.

Usage

```
npregtol.int(x, y, y.hat, side = 1, alpha = 0.05, P = 0.99,
            method = c("WILKS", "WALD", "HM"), upper = NULL,
            lower = NULL)
```

Arguments

x	A vector of values for the predictor variable. Currently, this function is only capable of handling a single predictor.
y	A vector of values for the response variable.
y.hat	A vector of fitted values extracted from a nonparametric smoothing routine.
side	Whether a 1-sided or 2-sided tolerance bound is required (determined by side = 1 or side = 2, respectively).
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of the population to be covered by the tolerance bound(s).
method	The method for determining which indices of the ordered residuals will be used for the tolerance bounds. "WILKS", "WALD", and "HM" are each described in nptol.int . However, since only one tolerance bound can actually be reported for this procedure, only the first tolerance bound will be returned. Note that this is not an issue when method = "WILKS" is used as it only produces one set of tolerance bounds.
upper	The upper bound of the data. When NULL, then the maximum of x is used.
lower	The lower bound of the data. When NULL, then the minimum of x is used.

Value

npregtol.int returns a data frame with items:

alpha	The specified significance level.
P	The proportion of the population covered by the tolerance bound(s).
x	The values of the predictor variable.
y	The values of the response variable.
y.hat	The predicted value of the response for the fitted nonparametric smoothing routine.
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

References

Young, D. S. (2013), Regression Tolerance Intervals, *Communications in Statistics - Simulation and Computation*, **42**, 2040–2055.

See Also

[loess](#), [nptol.int](#), [spline](#)

Examples

```
## 95%/95% 2-sided nonparametric regression tolerance bounds
## for a sample of size 50.

set.seed(100)
x <- runif(50, 5, 45)
f1 <- function(x, b1, b2) b1 + (0.49 - b1)*exp(-b2*(x - 8)) +
  rnorm(50, sd = 0.01)
y <- f1(x, 0.39, 0.11)
y.hat <- loess(y~x)$fit
out <- npregtol.int(x = x, y = y, y.hat = y.hat, side = 2,
  alpha = 0.05, P = 0.95, method = "WILKS")
out

plottol(out, x = x, y = y, y.hat = y.hat, side = "two",
  x.lab = "X", y.lab = "Y")
```

nptol.int

Nonparametric Tolerance Intervals

Description

Provides 1-sided or 2-sided nonparametric (i.e., distribution-free) tolerance intervals for any continuous data set.

Usage

```
nptol.int(x, alpha = 0.05, P = 0.99, side = 1,
  method = c("WILKS", "WALD", "HM", "YM"),
  upper = NULL, lower = NULL)
```

Arguments

x	A vector of data which no distributional assumptions are made. The data is only assumed to come from a continuous distribution.
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of the population to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by side = 1 or side = 2, respectively).

method	The method for determining which indices of the ordered observations will be used for the tolerance intervals. "WILKS" is the Wilks method, which produces tolerance bounds symmetric about the observed center of the data by using the beta distribution. "WALD" is the Wald method, which produces (possibly) multiple tolerance bounds for side = 2 (each having at least the specified confidence level), but is the same as method = "WILKS" for side = 1. "HM" is the Hahn-Meeker method, which is based on the binomial distribution, but the upper and lower bounds may exceed the minimum and maximum of the sample data. For side = 2, this method will yield two intervals if an odd number of observations are to be trimmed from each side. "YM" is the Young-Mathew method for performing interpolation or extrapolation based on the order statistics. See below for more information on this method.
upper	The upper bound of the data. When NULL, then the maximum of x is used. If method = "YM" and extrapolation is performed, then upper will be greater than the maximum.
lower	The lower bound of the data. When NULL, then the minimum of x is used. If method = "YM" and extrapolation is performed, then lower will be less than the maximum.

Details

For the Young-Mathew method, interpolation or extrapolation is performed. When side = 1, two intervals are given: one based on linear interpolation/extrapolation of order statistics (OS-Based) and one based on fractional order statistics (FOS-Based). When side = 2, only an interval based on linear interpolation/extrapolation of order statistics is given.

Value

nptol.int returns a data frame with items:

alpha	The specified significance level.
P	The proportion of the population covered by this tolerance interval.
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

References

- Bury, K. (1999), *Statistical Distributions in Engineering*, Cambridge University Press.
- Hahn, G. J. and Meeker, W. Q. (1991), *Statistical Intervals: A Guide for Practitioners*, Wiley-Interscience.
- Wald, A. (1943), An Extension of Wilks' Method for Setting Tolerance Limits, *The Annals of Mathematical Statistics*, **14**, 45–55.
- Wilks, S. S. (1941), Determination of Sample Sizes for Setting Tolerance Limits, *The Annals of Mathematical Statistics*, **12**, 91–96.

Young, D. S. and Mathew, T. (2014), Improved Nonparametric Tolerance Intervals Based on Interpolated and Extrapolated Order Statistics, *Journal of Nonparametric Statistics*, **26**, 415–432.

See Also

[distfree.est](#), [npregtol.int](#)

Examples

```
## 90%/95% 2-sided nonparametric tolerance intervals for a
## sample of size 20.

set.seed(100)
x <- rlogis(20, 5, 1)
out <- nptol.int(x = x, alpha = 0.10, P = 0.95, side = 1,
                method = "WILKS", upper = NULL, lower = NULL)
out

plottol(out, x, plot.type = "both", side = "two", x.lab = "X")
```

paretotol.int

Pareto (or Power Distribution) Tolerance Intervals

Description

Provides 1-sided or 2-sided tolerance intervals for data distributed according to either a Pareto distribution or a power distribution (i.e., the inverse Pareto distribution).

Usage

```
paretotol.int(x, alpha = 0.05, P = 0.99, side = 1,
              method = c("GPU", "DUN"), power.dist = FALSE)
```

Arguments

x	A vector of data which is distributed according to either a Pareto distribution or a power distribution.
alpha	The level chosen such that $1 - \alpha$ is the confidence level.
P	The proportion of the population to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by <code>side = 1</code> or <code>side = 2</code> , respectively).
method	The method for how the upper tolerance bound is approximated when transforming to utilize the relationship with the 2-parameter exponential distribution. "GPU" is the Guenther-Patil-Uppuluri method. "DUN" is the Dunsmore method, which was empirically shown to be an improvement for samples greater than or equal to 8. More information on these methods can be found in the "References".

`power.dist` If TRUE, then the data is considered to be from a power distribution, in which case the output gives tolerance intervals for the power distribution. The default is FALSE.

Details

Recall that if the random variable X is distributed according to a Pareto distribution, then the random variable $Y = \ln(X)$ is distributed according to a 2-parameter exponential distribution. Moreover, if the random variable W is distributed according to a power distribution, then the random variable $X = 1/W$ is distributed according to a Pareto distribution, which in turn means that the random variable $Y = \ln(1/W)$ is distributed according to a 2-parameter exponential distribution.

Value

`paretotol.int` returns a data frame with items:

<code>alpha</code>	The specified significance level.
<code>P</code>	The proportion of the population covered by this tolerance interval.
<code>1-sided.lower</code>	The 1-sided lower tolerance bound. This is given only if <code>side = 1</code> .
<code>1-sided.upper</code>	The 1-sided upper tolerance bound. This is given only if <code>side = 1</code> .
<code>2-sided.lower</code>	The 2-sided lower tolerance bound. This is given only if <code>side = 2</code> .
<code>2-sided.upper</code>	The 2-sided upper tolerance bound. This is given only if <code>side = 2</code> .

References

- Dunsmore, I. R. (1978), Some Approximations for Tolerance Factors for the Two Parameter Exponential Distribution, *Technometrics*, **20**, 317–318.
- Engelhardt, M. and Bain, L. J. (1978), Tolerance Limits and Confidence Limits on Reliability for the Two-Parameter Exponential Distribution, *Technometrics*, **20**, 37–39.
- Guenther, W. C., Patil, S. A., and Uppuluri, V. R. R. (1976), One-Sided β -Content Tolerance Factors for the Two Parameter Exponential Distribution, *Technometrics*, **18**, 333–340.
- Krishnamoorthy, K., Mathew, T., and Mukherjee, S. (2008), Normal-Based Methods for a Gamma Distribution: Prediction and Tolerance Intervals and Stress-Strength Reliability, *Technometrics*, **50**, 69–78.

See Also

[TwoParExponential](#), [exp2tol.int](#)

Examples

```
## 95%/99% 2-sided Pareto tolerance intervals
## for a sample of size 500.

set.seed(100)
x <- exp(r2exp(500, rate = 0.15, shift = 2))
out <- paretotol.int(x = x, alpha = 0.05, P = 0.99, side = 2,
```

```

                                method = "DUN", power.dist = FALSE)
out
plottol(out, x, plot.type = "both", side = "two",
        x.lab = "Pareto Data")

```

plottol

Plotting Capabilities for Tolerance Intervals

Description

Provides control charts and/or histograms for tolerance bounds on continuous data as well as tolerance ellipses for data distributed according to bivariate and trivariate normal distributions. Scatterplots with regression tolerance bounds and interval plots for ANOVA tolerance intervals may also be produced.

Usage

```

plottol(tol.out, x, y = NULL, y.hat = NULL,
        side = c("two", "upper", "lower"),
        plot.type = c("control", "hist", "both"),
        x.lab = NULL, y.lab = NULL, z.lab = NULL, ...)

```

Arguments

tol.out	Output from any continuous (including ANOVA) tolerance interval procedure or from a regression tolerance bound procedure.
x	Either data from a continuous distribution or the predictors for a regression model. If this is a design matrix for a linear regression model, then it must be in matrix form AND include a column of 1's if there is to be an intercept. Note that multiple predictors are only allowed if considering polynomial regression. If the output for tol.out concerns ANOVA tolerance intervals, then x must be a data frame.
y	The response vector for a regression setting. Leave as NULL if not doing regression tolerance bounds.
y.hat	The fitted values from a nonparametric smoothing routine if plotting nonparametric regression tolerance bounds. Otherwise, leave as NULL.
side	side = "two" produces plots for either the two-sided tolerance intervals or both one-sided tolerance intervals. This will be determined by the output in tol.out. side = "upper" produces plots showing the upper tolerance bounds. side = "lower" produces plots showing the lower tolerance bounds.
plot.type	plot.type = "control" produces a control chart of the data along with the tolerance bounds specified by side. plot.type = "hist" produces a histogram of the data along with the tolerance bounds specified by side. plot.type = "both" produces both the control chart and histogram. This argument is ignored when plotting regression data.

x.lab	Specify the label for the x-axis.
y.lab	Specify the label for the y-axis.
z.lab	Specify the label for the z-axis.
...	Additional arguments passed to the plotting function used for the control charts or regression scatterplots.

Value

plottol can return a control chart, histogram, or both for continuous data along with the calculated tolerance intervals. For regression data, plottol returns a scatterplot along with the regression tolerance bounds. For ANOVA output, plottol returns an interval plot for each factor.

References

Montgomery, D. C. (2005), *Introduction to Statistical Quality Control*, Fifth Edition, John Wiley & Sons, Inc.

Examples

```
## 90%/90% 1-sided Weibull tolerance intervals for a sample
## of size 150.

set.seed(100)
x <- rweibull(150, 3, 75)
out <- exttol.int(x = x, alpha = 0.15, P = 0.90,
                 dist = "Weibull")
out

plottol(out, x, plot.type = "both", side = "lower",
        x.lab = "Weibull Data")

## 90%/90% trivariate normal tolerance region.

set.seed(100)
x1 <- rnorm(100, 0, 0.2)
x2 <- rnorm(100, 0, 0.5)
x3 <- rnorm(100, 5, 1)
x <- cbind(x1, x2, x3)
mvtol.region(x = x, alpha = c(0.10, 0.05, 0.01),
             P = c(0.90, 0.95, 0.99), B = 1000)

out2 <- mvtol.region(x = x, alpha = 0.10, P = 0.90, B = 1000)
out2
plottol(out2, x)

## 95%/95% 2-sided linear regression tolerance bounds
## for a sample of size 100.

set.seed(100)
x <- runif(100, 0, 10)
```

```

y <- 20 + 5*x + rnorm(100, 0, 3)
out3 <- regtol.int(reg = lm(y ~ x), new.x = cbind(c(3, 6, 9)),
                 side = 2, alpha = 0.05, P = 0.95)
plottol(out3, x = cbind(1, x), y = y, side = "two", x.lab = "X",
        y.lab = "Y")

```

poistol.int

Poisson Tolerance Intervals

Description

Provides 1-sided or 2-sided tolerance intervals for Poisson random variables. From a statistical quality control perspective, these limits bound the number of occurrences (which follow a Poisson distribution) in a specified future time period.

Usage

```

poistol.int(x, n, m = NULL, alpha = 0.05, P = 0.99, side = 1,
           method = c("TAB", "LS", "SC", "CC", "VS", "RVS",
                    "FT", "CSC"))

```

Arguments

x	The number of occurrences of the event in time period n. Can be a vector of length n, in which case the sum of x is used.
n	The time period of the original measurements.
m	The specified future length of time. If m = NULL, then the tolerance limits will be constructed assuming n for the future length of time.
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of occurrences in future time lengths of size m to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by side = 1 or side = 2, respectively).
method	The method for calculating the lower and upper confidence bounds, which are used in the calculation of the tolerance bounds. The default method is "TAB", which is the tabular method and is usually preferred for a smaller number of occurrences. "LS" gives the large-sample (Wald) method, which is usually preferred when the number of occurrences is $x > 20$. "SC" gives the score method, which again is usually used when the number of occurrences is relatively large. "CC" gives a continuity-corrected version of the large-sample method. "VS" gives a variance-stabilized version of the large-sample method. "RVS" is a recentered version of the variance-stabilization method. "FT" is the Freeman-Tukey method. "CSC" is the continuity-corrected version of the score method. More information on these methods can be found in the "References".

Value

poistol.int returns a data frame with items:

alpha	The specified significance level.
P	The proportion of occurrences in future time periods of length m.
lambda.hat	The mean occurrence rate per unit time, calculated by x/n .
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

References

Barker, L. (2002), A Comparison of Nine Confidence Intervals for a Poisson Parameter When the Expected Number of Events Is ≤ 5 , *The American Statistician*, **56**, 85–89.

Freeman, M. F. and Tukey, J. W. (1950), Transformations Related to the Angular and the Square Root, *Annals of Mathematical Statistics*, **21**, 607–611.

Hahn, G. J. and Chandra, R. (1981), Tolerance Intervals for Poisson and Binomial Variables, *Journal of Quality Technology*, **13**, 100–110.

See Also

[Poisson](#), [umatol.int](#)

Examples

```
## 95%/90% 1-sided Poisson tolerance limits for future
## occurrences in a period of length 3. All seven methods
## are presented for comparison.

poistol.int(x = 45, n = 9, m = 3, alpha = 0.05, P = 0.90,
            side = 1, method = "TAB")
poistol.int(x = 45, n = 9, m = 3, alpha = 0.05, P = 0.90,
            side = 1, method = "LS")
poistol.int(x = 45, n = 9, m = 3, alpha = 0.05, P = 0.90,
            side = 1, method = "SC")
poistol.int(x = 45, n = 9, m = 3, alpha = 0.05, P = 0.90,
            side = 1, method = "CC")
poistol.int(x = 45, n = 9, m = 3, alpha = 0.05, P = 0.90,
            side = 1, method = "VS")
poistol.int(x = 45, n = 9, m = 3, alpha = 0.05, P = 0.90,
            side = 1, method = "RVS")
poistol.int(x = 45, n = 9, m = 3, alpha = 0.05, P = 0.90,
            side = 1, method = "FT")
poistol.int(x = 45, n = 9, m = 3, alpha = 0.05, P = 0.90,
            side = 1, method = "CSC")
```

```
## 95%/90% 2-sided Poisson tolerance intervals for future
## occurrences in a period of length 15. All seven methods
## are presented for comparison.
```

```
poistol.int(x = 45, n = 9, m = 15, alpha = 0.05, P = 0.90,
            side = 2, method = "TAB")
poistol.int(x = 45, n = 9, m = 15, alpha = 0.05, P = 0.90,
            side = 2, method = "LS")
poistol.int(x = 45, n = 9, m = 15, alpha = 0.05, P = 0.90,
            side = 2, method = "SC")
poistol.int(x = 45, n = 9, m = 15, alpha = 0.05, P = 0.90,
            side = 2, method = "CC")
poistol.int(x = 45, n = 9, m = 15, alpha = 0.05, P = 0.90,
            side = 2, method = "VS")
poistol.int(x = 45, n = 9, m = 15, alpha = 0.05, P = 0.90,
            side = 2, method = "RVS")
poistol.int(x = 45, n = 9, m = 15, alpha = 0.05, P = 0.90,
            side = 2, method = "FT")
poistol.int(x = 45, n = 9, m = 15, alpha = 0.05, P = 0.90,
            side = 2, method = "CSC")
```

regtol.int

(Multiple) Linear Regression Tolerance Bounds

Description

Provides 1-sided or 2-sided (multiple) linear regression tolerance bounds. It is also possible to fit a regression through the origin model.

Usage

```
regtol.int(reg, new.x = NULL, side = 1, alpha = 0.05, P = 0.99)
```

Arguments

reg	An object of class <code>lm</code> (i.e., the results from a linear regression routine).
new.x	Any new levels of the predictor(s) for which to report the tolerance bounds. The number of columns must equal the number of predictors appearing on the right-hand side of the model in <code>reg</code> .
side	Whether a 1-sided or 2-sided tolerance bound is required (determined by <code>side = 1</code> or <code>side = 2</code> , respectively).
alpha	The level chosen such that $1-\alpha$ is the confidence level.
P	The proportion of the population to be covered by the tolerance bound(s).

Value

regtol.int returns a data frame with items:

alpha	The specified significance level.
P	The proportion of the population covered by the tolerance bound(s).
y	The value of the response given on the left-hand side of the model in reg.
y.hat	The predicted value of the response for the fitted linear regression model. This data frame is sorted by this value.
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

References

Wallis, W. A. (1951), Tolerance Intervals for Linear Regression, in *Second Berkeley Symposium on Mathematical Statistics and Probability*, ed. J. Neyman, Berkeley: University of CA Press, 43–51.

Young, D. S. (2013), Regression Tolerance Intervals, *Communications in Statistics - Simulation and Computation*, **42**, 2040–2055.

See Also

[lm](#)

Examples

```
## 95%/95% 2-sided linear regression tolerance bounds
## for a sample of size 100.

set.seed(100)
x <- runif(100, 0, 10)
y <- 20 + 5*x + rnorm(100, 0, 3)
out <- regtol.int(reg = lm(y ~ x), new.x = cbind(c(3, 6, 9)),
                 side = 2, alpha = 0.05, P = 0.95)

out

plottol(out, x = cbind(1, x), y = y, side = "two", x.lab = "X",
        y.lab = "Y")
```

TwoParExponential *The 2-Parameter Exponential Distribution*

Description

Density, distribution function, quantile function, and random generation for the 2-parameter exponential distribution with rate equal to rate and shift equal to shift.

Usage

```
d2exp(x, rate = 1, shift = 0, log = FALSE)
p2exp(q, rate = 1, shift = 0, lower.tail = TRUE, log.p = FALSE)
q2exp(p, rate = 1, shift = 0, lower.tail = TRUE, log.p = FALSE)
r2exp(n, rate = 1, shift = 0)
```

Arguments

x, q	Vector of quantiles.
p	Vector of probabilities.
n	The number of observations. If length>1, then the length is taken to be the number required.
rate	Vector of rates.
shift	Vector of shifts.
log, log.p	Logical vectors. If TRUE, then probabilities are given as log(p).
lower.tail	Logical vector. If TRUE, then probabilities are $P[X \leq x]$, else $P[X > x]$.

Details

If rate or shift are not specified, then they assume the default values of 1 and 0, respectively.

The 2-parameter exponential distribution has density

$$f(x) = \frac{1}{\beta} e^{-(x-\mu)/\beta}$$

where $x \geq \mu$, μ is the shift parameter, and $\beta > 0$ is the scale parameter.

Value

d2exp gives the density, p2exp gives the distribution function, q2exp gives the quantile function, and r2exp generates random deviates.

See Also

[runif](#) and [.Random.seed](#) about random number generation.

Examples

```
## Randomly generated data from the 2-parameter exponential
## distribution.

set.seed(100)
x <- r2exp(n = 500, rate = 3, shift = -10)
hist(x, main = "Randomly Generated Data", prob = TRUE)

x.1 = sort(x)
y <- d2exp(x = x.1, rate = 3, shift = -10)
lines(x.1, y, col = 2, lwd = 2)

plot(x.1, p2exp(q = x.1, rate = 3, shift = -10), type = "l",
      xlab = "x", ylab = "Cumulative Probabilities")

q2exp(p = 0.20, rate = 3, shift = -10, lower.tail = FALSE)
q2exp(p = 0.80, rate = 3, shift = -10)
```

umatol.int

Uniformly Most Accurate Upper Tolerance Limits for Certain Discrete Distributions

Description

Provides uniformly most accurate upper tolerance limits for the binomial, negative binomial, and Poisson distributions.

Usage

```
umatol.int(x, n = NULL, dist = c("Bin", "NegBin", "Pois"), N,
           alpha = 0.05, P = 0.99)
```

Arguments

x	A vector of data which is distributed according to one of the binomial, negative binomial, or Poisson distributions. If the length of x is 1, then it is assumed that this number is the sum of iid values from the assumed distribution.
n	The sample size of the data. If null, then n is calculated as the length of x.
dist	The distribution for the data given by x. The options are "Bin" for the binomial distribution, "NegBin" for the negative binomial distribution, and "Pois" for the Poisson distribution.
N	Must be specified for the binomial and negative binomial distributions. If dist = "Bin", then N is the number of Bernoulli trials and must be a positive integer. If dist = "NegBin", then N is the total number of successful trials (or dispersion parameter) and must be strictly positive.
alpha	The level chosen such that 1-alpha is the confidence level.
P	The proportion of the population to be covered by this tolerance interval.

Value

umatol.int returns a data frame with items:

alpha	The specified significance level.
P	The proportion of the population covered by this tolerance interval.
p.hat	The maximum likelihood estimate for the probability of success in each trial; reported if dist = "Bin".
nu.hat	The maximum likelihood estimate for the probability of success in each trial; reported if dist = "NegBin".
lambda.hat	The maximum likelihood estimate for the rate of success; reported if dist = "Pois".
1-sided.upper	The 1-sided upper tolerance limit.

References

Zacks, S. (1970), Uniformly Most Accurate Tolerance Limits for Monotone Likelihood Ratio Families of Discrete Distributions, *Journal of the American Statistical Association*, **65**, 307–316.

See Also

[Binomial](#), [NegBinomial](#), [Poisson](#)

Examples

```
## Examples from Zacks (1970).

umatol.int(25, n = 4, dist = "Bin", N = 10, alpha = 0.10,
           P = 0.95)
umatol.int(13, n = 10, dist = "NegBin", N = 2, alpha = 0.10,
           P = 0.95)
umatol.int(37, n = 10, dist = "Pois", alpha = 0.10, P = 0.95)
```

uniftol.int

Uniform Tolerance Intervals

Description

Provides 1-sided or 2-sided tolerance intervals for data distributed according to a uniform distribution.

Usage

```
uniftol.int(x, alpha = 0.05, P = 0.99, upper = NULL,
           lower = NULL, side = 1)
```

Arguments

x	A vector of data which is distributed according to a uniform distribution.
alpha	The level chosen such that $1-\alpha$ is the confidence level.
P	The proportion of the population to be covered by this tolerance interval.
upper	The upper bound of the data. When NULL, then the maximum of x is used.
lower	The lower bound of the data. When NULL, then the minimum of x is used.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by side = 1 or side = 2, respectively).

Value

uniftol.int returns a data frame with items:

alpha	The specified significance level.
P	The proportion of the population covered by this tolerance interval.
1-sided.lower	The 1-sided lower tolerance bound. This is given only if side = 1.
1-sided.upper	The 1-sided upper tolerance bound. This is given only if side = 1.
2-sided.lower	The 2-sided lower tolerance bound. This is given only if side = 2.
2-sided.upper	The 2-sided upper tolerance bound. This is given only if side = 2.

References

Faulkenberry, G. D. and Weeks, D. L. (1968), Sample Size Determination for Tolerance Limits, *Technometrics*, **10**, 343–348.

Examples

```
## 90%/90% 1-sided uniform tolerance intervals for a sample
## of size 50 with a known lower bound of 0.

set.seed(100)
x <- runif(50, 0, 50)
out <- uniftol.int(x = x, alpha = 0.10, P = 0.90, lower = 0,
                  side = 1)

out

plottol(out, x, plot.type = "hist", side = "two",
        x.lab = "Uniform Data")
```

Description

Density (mass), distribution function, quantile function, and random generation for the Zipf, Zipf-Mandelbrot, and zeta distributions.

Usage

```

dzipfman(x, s, b = NULL, N = NULL, log = FALSE)
pzipfman(q, s, b = NULL, N = NULL, lower.tail = TRUE,
         log.p = FALSE)
qzipfman(p, s, b = NULL, N = NULL, lower.tail = TRUE,
         log.p = FALSE)
rzipfman(n, s, b = NULL, N = NULL)

```

Arguments

<code>x, q</code>	Vector of quantiles.
<code>p</code>	Vector of probabilities.
<code>n</code>	The number of observations. If <code>length>1</code> , then the length is taken to be the number required.
<code>s, b</code>	The shape parameters, both of which must be greater than 0. <code>b</code> must be specified for Zipf-Mandelbrot distributions.
<code>N</code>	The number of categories, which must be integer-valued for Zipf and Zipf-Mandelbrot distributions. For a zeta distribution, <code>N = Inf</code> must be used.
<code>log, log.p</code>	Logical vectors. If TRUE, then the probabilities are given as $\log(p)$.
<code>lower.tail</code>	Logical vector. If TRUE, then probabilities are $P[X \leq x]$, else $P[X > x]$.

Details

The Zipf-Mandelbrot distribution has mass

$$p(x) = \frac{(x + \delta)^{-\lambda}}{\sum_{i=1}^N (i + \delta)^{-\lambda}},$$

where $x = 1, \dots, N$, $\lambda, \delta > 0$ are shape parameters, and N is the number of distinct categories. The Zipf distribution is just a special case of the Zipf-Mandelbrot distribution where the second shape parameter $b=0$. The zeta distribution has mass

$$p(x) = \frac{x^{-\lambda}}{\zeta(\lambda)},$$

where $x = 1, 2, \dots$, $\lambda > 1$ is the shape parameter, and $\zeta()$ is the Riemann zeta function given by:

$$\zeta(t) = \sum_{i=1}^{\infty} \frac{1}{i^t} < \infty.$$

Note that the zeta distribution is just a special case of the Zipf distribution where $s > 1$ and N goes to infinity.

Value

`dzipfman` gives the density (mass), `pzipfman` gives the distribution function, `qzipfman` gives the quantile function, and `rzipfman` generates random deviates for the specified distribution.

Note

These functions may be updated in a future version of the package so as to allow greater flexibility with the inputs.

References

- Mandelbrot, B. B. (1965), Information Theory and Psycholinguistics. In B. B. Wolman and E. Nagel, editors. *Scientific Psychology*, Basic Books.
- Young, D. S. (2013), Approximate Tolerance Limits for Zipf-Mandelbrot Distributions, *Physica A: Statistical Mechanics and its Applications*, **392**, 1702–1711.
- Zipf, G. K. (1949), *Human Behavior and the Principle of Least Effort*, Hafner.
- Zornig, P. and Altmann, G. (1995), Unified Representation of Zipf Distributions, *Computational Statistics and Data Analysis*, **19**, 461–473.

See Also

[runif](#) and [.Random.seed](#) about random number generation.

Examples

```
## Randomly generated data from the Zipf distribution.

set.seed(100)
x <- rzipfman(n = 150, s = 2, N = 100)
hist(x, main = "Randomly Generated Data", prob = TRUE)

x.1 <- sort(x)
y <- dzipfman(x = x.1, s = 2, N = 100)
lines(x.1, y, col = 2, lwd = 2)

plot(x.1, pzipfman(q = x.1, s = 2, N = 100), type = "l",
      xlab = "x", ylab = "Cumulative Probabilities")

qzipfman(p = 0.20, s = 2, N = 100, lower.tail = FALSE)
qzipfman(p = 0.80, s = 2, N = 100)

## Randomly generated data from the Zipf-Mandelbrot distribution.

set.seed(100)
x <- rzipfman(n = 150, s = 2, b = 3, N = 100)
hist(x, main = "Randomly Generated Data", prob = TRUE)
```

```

x.1 <- sort(x)
y <- dzipfman(x = x.1, s = 2, b = 3, N = 100)
lines(x.1, y, col = 2, lwd = 2)

plot(x.1, pzipfman(q = x.1, s = 2, b = 3, N = 100), type = "l",
      xlab = "x", ylab = "Cumulative Probabilities")

qzipfman(p = 0.20, s = 2, b = 3, N = 100, lower.tail = FALSE)
qzipfman(p = 0.80, s = 2, b = 3, N = 100)

## Randomly generated data from the zeta distribution.

set.seed(100)
x <- rzipfman(n = 100, s = 1.3, N = Inf)
hist(x, main = "Randomly Generated Data", prob = TRUE)

x.1 <- sort(x)
y <- dzipfman(x = x.1, s = 1.3, N = Inf)
lines(x.1, y, col = 2, lwd = 2)

plot(x.1, pzipfman(q = x.1, s = 1.3, N = Inf), type = "l",
      xlab = "x", ylab = "Cumulative Probabilities")

qzipfman(p = 0.20, s = 1.3, lower.tail = FALSE, N = Inf)
qzipfman(p = 0.80, s = 1.3, N = Inf)

```

zipftol.int

Zipf-Mandelbrot Tolerance Intervals

Description

Provides 1-sided or 2-sided tolerance intervals for data distributed according to Zipf, Zipf-Mandelbrot, and zeta distributions.

Usage

```

zipftol.int(x, m = NULL, N = NULL, alpha = 0.05, P = 0.99,
            side = 1, s = 1, b = 1, dist = c("Zipf",
            "Zipf-Man", "Zeta"), ties = FALSE, ...)

```

Arguments

x	A vector of raw data or a table of counts which is distributed according to a Zipf, Zipf-Mandelbrot, or zeta distribution. Do not supply a vector of counts!
m	The number of observations in a future sample for which the tolerance limits will be calculated. By default, m = NULL and, thus, m will be set equal to the original sample size.

N	The number of categories when <code>dist = "Zipf"</code> or <code>dist = "Zipf-Man"</code> . This is not used when <code>dist = "Zeta"</code> . If <code>N = NULL</code> , then <code>N</code> is estimated based on the number of categories observed in the data.
alpha	The level chosen such that $1-\alpha$ is the confidence level.
P	The proportion of the population to be covered by this tolerance interval.
side	Whether a 1-sided or 2-sided tolerance interval is required (determined by <code>side = 1</code> or <code>side = 2</code> , respectively).
s	The initial value to estimate the shape parameter in the <code>zm.ll</code> function.
b	The initial value to estimate the second shape parameter in the <code>zm.ll</code> function when <code>dist = "Zipf-Man"</code> .
dist	Options are <code>dist = "Zipf"</code> , <code>dist = "Zipf-Man"</code> , or <code>dist = "Zeta"</code> if the data is distributed according to the Zipf, Zipf-Mandelbrot, or zeta distribution, respectively.
ties	How to handle if there are other categories with the same frequency as the category at the estimated tolerance limit. The default is <code>FALSE</code> , which does no correction. If <code>TRUE</code> , then the highest ranked (i.e., lowest number) of the tied categories is selected for the lower limit and the lowest ranked (i.e., highest number) of the tied categories is selected for the upper limit.
...	Additional arguments passed to the <code>zm.ll</code> function, which is used for maximum likelihood estimation.

Details

Zipf-Mandelbrot models are commonly used to model phenomena where the frequencies of categorical data are approximately inversely proportional to its rank in the frequency table. Zipf-Mandelbrot distributions are heavily right-skewed distributions with a (relatively) large mass placed on the first category. For most practical applications, one will typically be interested in 1-sided upper bounds.

Value

`zipftol.int` returns a data frame with the following items:

alpha	The specified significance level.
P	The proportion of the population covered by this tolerance interval.
s.hat	MLE for the shape parameter <code>s</code> .
b.hat	MLE for the shape parameter <code>b</code> when <code>dist = "Zipf-Man"</code> .
1-sided.lower	The 1-sided lower tolerance bound. This is given only if <code>side = 1</code> .
1-sided.upper	The 1-sided upper tolerance bound. This is given only if <code>side = 1</code> .
2-sided.lower	The 2-sided lower tolerance bound. This is given only if <code>side = 2</code> .
2-sided.upper	The 2-sided upper tolerance bound. This is given only if <code>side = 2</code> .

Note

This function may be updated in a future version of the package so as to allow greater flexibility with the inputs.

References

- Mandelbrot, B. B. (1965), Information Theory and Psycholinguistics. In B. B. Wolman and E. Nagel, editors. *Scientific Psychology*, Basic Books.
- Young, D. S. (2013), Approximate Tolerance Limits for Zipf-Mandelbrot Distributions, *Physica A: Statistical Mechanics and its Applications*, **392**, 1702–1711.
- Zipf, G. K. (1949), *Human Behavior and the Principle of Least Effort*, Hafner.
- Zornig, P. and Altmann, G. (1995), Unified Representation of Zipf Distributions, *Computational Statistics and Data Analysis*, **19**, 461–473.

See Also

[ZipfMandelbrot](#), [zm.ll](#)

Examples

```
## 95%/99% 1-sided tolerance intervals for the Zipf,  
## Zipf-Mandelbrot, and zeta distributions.  
  
set.seed(100)  
  
s <- 2  
b <- 5  
N <- 50  
  
zipf.data <- rzipfman(n = 150, s = s, N = N)  
zipfman.data <- rzipfman(n = 150, s = s, b = b, N = N)  
zeta.data <- rzipfman(n = 150, s = s, N = Inf)  
  
out.zipf <- zipftol.int(zipf.data, dist = "Zipf")  
out.zipfman <- zipftol.int(zipfman.data, dist = "Zipf-Man")  
out.zeta <- zipftol.int(zeta.data, N = Inf, dist = "Zeta")  
  
out.zipf  
out.zipfman  
out.zeta
```

zm.ll

Maximum Likelihood Estimation for Zipf-Mandelbrot Models

Description

Performs maximum likelihood estimation for the parameters of the Zipf, Zipf-Mandelbrot, and zeta distributions.

Usage

```
zm.ll(x, N = NULL, s = 1, b = 1, dist = c("Zipf", "Zipf-Man",  
    "Zeta"), ...)
```


Arguments

x	A vector of raw data or a table of counts which is distributed according to a Zipf, Zipf-Mandelbrot, or zeta distribution. Do not supply a vector of counts!
N	The number of categories when <code>dist = "Zipf"</code> or <code>dist = "Zipf-Man"</code> . This is not used when <code>dist = "Zeta"</code> . If <code>N = NULL</code> , then <code>N</code> is estimated based on the number of categories observed in the data.
s	The initial value to estimate the shape parameter, which is set to 1 by default. If a poor initial value is specified, then a <code>WARNING</code> message is returned.
b	The initial value to estimate the second shape parameter when <code>dist = "Zipf-Man"</code> , which is set to 1 by default. If a poor initial value is specified, then a <code>WARNING</code> message is returned.
dist	Options are <code>dist = "Zipf"</code> , <code>dist = "Zipf-Man"</code> , or <code>dist = "Zeta"</code> if the data is distributed according to the Zipf, Zipf-Mandelbrot, or zeta distribution, respectively.
...	Additional arguments passed to the <code>mle</code> function.

Details

Zipf-Mandelbrot models are commonly used to model phenomena where the frequencies of categorical data are approximately inversely proportional to its rank in the frequency table.

Value

See the help file for `mle` to see how the output is structured.

Note

This function may be updated in a future version of the package so as to allow greater flexibility with the inputs.

References

Mandelbrot, B. B. (1965), Information Theory and Psycholinguistics. In B. B. Wolman and E. Nagel, editors. *Scientific Psychology*, Basic Books.

Zipf, G. K. (1949), *Human Behavior and the Principle of Least Effort*, Hafner.

Zornig, P. and Altmann, G. (1995), Unified Representation of Zipf Distributions, *Computational Statistics and Data Analysis*, **19**, 461–473.

See Also

[mle](#), [ZipfMandelbrot](#)

Examples

```
## Maximum likelihood estimation for randomly generated data
## from the Zipf, Zipf-Mandelbrot, and zeta distributions.

set.seed(100)

s <- 2
b <- 5
N <- 50

zipf.data <- rzipfman(n = 500, s = s, N = N)
out.zipf <- zm.ll(zipf.data, N = N, dist = "Zipf")
coef(out.zipf)
vcov(out.zipf)

zipfman.data <- rzipfman(n = 500, s = s, b = b, N = N)
out.zipfman <- zm.ll(zipfman.data, N = N, dist = "Zipf-Man")
coef(out.zipfman)
diag(vcov(out.zipfman))

zeta.data <- rzipfman(n = 200, s = s, N = Inf)
out.zeta <- zm.ll(zeta.data, N = Inf, dist = "Zeta")
coef(out.zeta)
vcov(out.zeta)
```

Index

*Topic file

acc.samp, 3
anovatul.int, 5
bintol.int, 7
bonftol.int, 9
cautol.int, 11
diffnormtol.int, 12
DiffProp, 14
distfree.est, 15
exp2tol.int, 17
exptol.int, 18
exttol.int, 20
F1, 22
fidbintol.int, 23
fidnegbintol.int, 25
fidpoistol.int, 27
gamtol.int, 29
hypertol.int, 31
K.factor, 32
K.table, 34
laptol.int, 36
logistol.int, 37
mvregtol.region, 38
mvtol.region, 40
negbintol.int, 42
NegHypergeometric, 44
neghypertol.int, 46
nlregtol.int, 47
normtol.int, 49
np.order, 51
npregtol.int, 52
nptol.int, 54
paretotol.int, 56
plottol, 58
poistol.int, 60
regtol.int, 62
TwoParExponential, 64
umatol.int, 65
uniftol.int, 66
ZipfMandelbrot, 68
zipftol.int, 70
zm.ll, 72
.Random.seed, 15, 45, 64, 69
acc.samp, 3, 32, 47
anova, 6
anovatul.int, 5
Binomial, 8, 66
bintol.int, 7
bonftol.int, 9
Cauchy, 11
cautol.int, 11
confint, 3
d2exp (TwoParExponential), 64
ddiffprop (DiffProp), 14
diffnormtol.int, 12
DiffProp, 14, 22
distfree.est, 15, 56
dnhyper (NegHypergeometric), 44
dzipfman (ZipfMandelbrot), 68
exp2tol.int, 17, 57
Exponential, 19
exptol.int, 18
exttol.int, 20
F1, 22
fidbintol.int, 23, 26, 28
fidnegbintol.int, 24, 25, 28
fidpoistol.int, 24, 26, 27
GammaDist, 30
gamtol.int, 29
Hypergeometric, 5, 32
hypertol.int, 31
integrate, 22, 34

K. factor, [6](#), [13](#), [30](#), [32](#), [35](#), [51](#)
K. table, [34](#), [34](#)

`laptol.int`, [36](#)
`lm`, [6](#), [63](#)
`loess`, [54](#)
`Logistic`, [38](#)
`logistol.int`, [37](#)

`mle`, [73](#)
`mvregtol.region`, [38](#)
`mvtol.region`, [40](#)

`NegBinomial`, [43](#), [66](#)
`negbintol.int`, [42](#)
`NegHypergeometric`, [44](#), [47](#)
`neghypertol.int`, [46](#)
`nlregtol.int`, [47](#)
`nls`, [48](#)
`Normal`, [13](#), [51](#)
`normtol.int`, [6](#), [13](#), [34](#), [49](#)
`np.order`, [51](#)
`npregtol.int`, [52](#), [56](#)
`nptol.int`, [16](#), [52–54](#), [54](#)

`p2exp (TwoParExponential)`, [64](#)
`paretotol.int`, [56](#)
`pdiffprop (DiffProp)`, [14](#)
`plottol`, [58](#)
`pnhyper (NegHypergeometric)`, [44](#)
`Poisson`, [61](#), [66](#)
`poistol.int`, [60](#)
`pzipfman (ZipfMandelbrot)`, [68](#)

`q2exp (TwoParExponential)`, [64](#)
`qdiffprop (DiffProp)`, [14](#)
`qnhyper (NegHypergeometric)`, [44](#)
`qzipfman (ZipfMandelbrot)`, [68](#)

`r2exp (TwoParExponential)`, [64](#)
`rdiffprop (DiffProp)`, [14](#)
`regtol.int`, [62](#)
`rnhyper (NegHypergeometric)`, [44](#)
`runif`, [15](#), [45](#), [64](#), [69](#)
`rzipfman (ZipfMandelbrot)`, [68](#)

`spline`, [54](#)

`TDist`, [33](#), [34](#)
`tolerance (tolerance-package)`, [2](#)

`tolerance-package`, [2](#)
`TwoParExponential`, [18](#), [57](#), [64](#)

`umatol.int`, [8](#), [43](#), [61](#), [65](#)
`uniftol.int`, [66](#)

`Weibull`, [21](#)

`ZipfMandelbrot`, [68](#), [72](#), [73](#)
`zipftol.int`, [70](#)
`zm.ll`, [72](#), [72](#)