

Package ‘DPQ’

October 18, 2019

Title Density, Probability, Quantile ('DPQ') Computations

Version 0.3-5

Date 2019-10-18

Description Computations for approximations and alternatives for the 'DPQ' (Density (pdf), Probability (cdf) and Quantile) functions for probability distributions in R.

Primary focus is on (central and non-central) beta, gamma and related distributions such as the chi-squared, F, and t.

--
This is for the use of researchers in these numerical approximation implementations, notably for my own use in order to improve R's own pbeta(), qgamma(), ..., etc: {``dpq''-functions}.

-- We plan to complement with 'DPQmpfr' to be suggested later.

Depends R (>= 3.5.0)

Imports stats, graphics, methods, utils, sfsmisc

Suggests Rmpfr, Matrix, mgcv, scatterplot3d, akima

SuggestsNote Matrix only for its ``test-tools-1.R''; mgcv,scatt.,akima:
some tests/

License GPL (>= 2)

Encoding UTF-8

Author Martin Maechler [aut, cre] <<https://orcid.org/0000-0002-8685-9910>>,
Morten Welinder [ctb] (pgamma C code, see PR#7307, Jan. 2005),
Wolfgang Viechtbauer [ctb] (dtWV(), 2002),
Ross Ihaka [ctb] (src/qchisq_appr.c),
Marius Hofert [ctb] (lsum(), lsum()),
R-core [ctb] (src/{dpq.h, algdiv.c, pnchisq.c}),
R Foundation [cph] (src/qchisq-appr.c)

Maintainer Martin Maechler <maechler@stat.math.ethz.ch>

Repository R-Forge

Repository/R-Forge/Project specfun

Repository/R-Forge/Revision 131

Repository/R-Forge/DateTimeStamp 2019-10-18 13:25:37

Date/Publication 2019-10-18 13:25:37

R topics documented:

DPQ-package	2
algdiv	5
b_chi	6
dchisqApprox	9
dgamma-utils	10
dgamma.R	11
dnt	12
dtWV	14
format01prec	16
lbeta	17
lgamma1p	19
log1mexp	20
log1pmx	21
logcf	22
logspace.add	23
lssum	24
lsum	25
newton	26
numer-utils	29
pbetaRv1	31
pl2curves	32
pnbeta	33
pnchi1sq	35
pnchisqAppr	38
pnchisqWienergerm	42
pnt	44
ppoisson	46
qbetaAppr	48
qchisqAppr	50
qgammaAppr	51
qnchisqAppr	53
qnormAppr	55
qtAppr	57
r_pois	58
Index	60

Description

Computations for approximations and alternatives for the 'DPQ' (Density (pdf), Probability (cdf) and Quantile) functions for probability distributions in R. Primary focus is on (central and non-central) beta, gamma and related distributions such as the chi-squared, F, and t. – This is for the use of researchers in these numerical approximation implementations, notably for my own use in order to improve R's own pbeta(), qgamma(), ..., etc: '"dpq"-functions. – We plan to complement with 'DPQmpfr' to be suggested later.

Details

The DESCRIPTION file:

Package:	DPQ
Title:	Density, Probability, Quantile ('DPQ') Computations
Version:	0.3-5
Date:	2019-10-18
Authors@R:	c(person("Martin", "Maechler", role=c("aut", "cre"), email="maechler@stat.math.ethz.ch"))
Description:	Computations for approximations and alternatives for the 'DPQ' (Density (pdf), Probability, Quantile) functions.
Depends:	R (>= 3.5.0)
Imports:	stats, graphics, methods, utils, sfsmisc
Suggests:	Rmpfr, Matrix, mgcv, scatterplot3d, akima
SuggestsNote:	Matrix only for its "test-tools-1.R"; mgcv,scatt.,akima: some tests/
License:	GPL (>= 2)
Encoding:	UTF-8
Author:	Martin Maechler [aut, cre] (< https://orcid.org/0000-0002-8685-9910 >), Morten Wicksell [ctb]
Maintainer:	Martin Maechler <maechler@stat.math.ethz.ch>
Repository:	R-Forge
Repository/R-Forge/Project:	specfun
Repository/R-Forge/Revision:	131
Repository/R-Forge/DateTimeStamp:	2019-10-18 13:25:37
Date/Publication:	2019-10-18 13:25:37

Index of help topics:

DPQ-package	Density, Probability, Quantile ('DPQ') Computations
M_LN2	Numerical Utilities - Functions, Constants
algdiv	Compute log(gamma(b))/gamma(a+b)) when b >= 8
b_chi	Compute E[chi_nu]/sqrt(nu) useful for t- and chi-Distributions
bd0	Utility Functions for 'dgamma()' - Pure R Versions
dgamma.R	Gamma Density Function Alternatives
dnchisqR	Approximations of the (Noncentral) Chi-Squared Density
dntJKBf1	Non-central t-Distribution Density - Algorithms and Approximations
dtWV	Noncentral t Distribution Density by W.V.
format01prec	Format Numbers in [0,1] with "Precise" Result
lbetaM	(Log) Beta Approximations
lgamma1p	Accurate 'log(gamma(a+1))'
log1mexp	Compute f(a) = log(1 - exp(-a)) Numerically Optimally
log1pmx	Accurate 'log(1+x) - x'
logcf	Continued Fraction Approximation of Log-Related Series
logspace.add	Logspace Arithmetix - Addition and Subtraction
lssum	Compute Logarithm of a Sum with Signed Large Summands
lsum	Properly Compute the Logarithm of a Sum

newton	Simple R level Newton Algorithm, Mostly for Didactical Reasons
pbetaRv1	Pure R Implementation of Old pbeta()
pchisqV	Wienergerm Approximations to (Non-Central) Chi-squared Probabilities
pl2curves	Plot 2 Noncentral Distribution Curves for Visual Comparison
pnbetaAppr2	Noncentral Beta Probabilities
pnchi1sq	(Probabilities of Non-Central Chi-squared Distribution for Special Cases
pnchisq	(Approximate) Probabilities of Non-Central Chi-squared Distribution
pntR	Non-central t Probability Distribution - Algorithms and Approximations
ppoisErr	Direct Computation of 'ppois()' Poisson Distribution Probabilities
qbetaAppr	Compute (Approximate) Quantiles of the Beta Distribution
qchisqAppr	Compute Approximate Quantiles of the Chi-Squared Distribution
qgammaAppr	Compute (Approximate) Quantiles of the Gamma Distribution
qnchisqAppr	Compute Approximate Quantiles of Noncentral Chi-Squared Distribution
qnormAppr	Approximations to 'qnorm()', i.e., z_alpha
qtAppr	Compute Approximate Quantiles of Non-Central t Distribution
r_pois	Compute Relative Size of i-th term of Poisson Distribution Series

Further information is available in the following vignettes:

Noncentral-Chisq	Noncentral Chi-Squared Probabilities – Algorithms in R (source)
comp-beta	Computing Beta(a,b) for Large Arguments (source)

An important goal is to investigate diverse algorithms and approximations of R's own density ($d*$ ()), probability ($p*$ ()), and quantile ($q*$ ()) functions, notably in "border" cases where the traditional published algorithms have shown to be suboptimal, not quite accurate, or even useless.

Examples are border cases of the beta distribution, or **non-central** distributions such as the non-central chi-squared and t-distributions.

Author(s)

Principal author and maintainer: Martin Maechler <maechler@stat.math.ethz.ch>

See Also

The package **DPQmpfr** (not yet on CRAN), which builds on this package and on **Rmpfr**.

Examples

```
## Show problem in R's non-central t-distrib. density:
example(dnt)
```

algdiv

Compute log(gamma(b)/gamma(a+b)) when b >= 8

Description

Computes

$$\text{algdiv}(a, b) := \log \frac{\Gamma(b)}{\Gamma(a + b)} = \log \Gamma(b) - \log \Gamma(a + b) = \text{lgamma}(b) - \text{lgamma}(a+b)$$

in a numerically stable way.

This is an auxiliary function in R's (TOMS 708) implementation of [pbeta\(\)](#), aka the incomplete beta function ratio.

Usage

```
algdiv(a, b)
```

Arguments

a, b	numeric vectors which will be recycled to the same length.
------	--

Details

Note that this is also useful to compute the Beta function

$$B(a, b) = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a + b)}.$$

Clearly,

$$\log B(a, b) = \log \Gamma(a) + \text{algdiv}(a, b) = \log \Gamma(a) - \log Q_{ab}(a, b)$$

In our `./tests/qbeta-dist.R` we look into computing `log(p*Beta(p,q))` accurately for `p << q`

We are proposing a nice solution there.

How is this related to `algdiv()` ?

Value

a numeric vector of length `max(length(a), length(b))` (if neither is of length 0, in which case the result has length 0 as well).

Author(s)

Didonato, A. and Morris, A., Jr, (1992); `algdiv()`'s C version from the R sources, authored by the R core team; C and R interface: Martin Maechler

References

Didonato, A. and Morris, A., Jr, (1992) Algorithm 708: Significant digit computation of the incomplete beta function ratios, *ACM Transactions on Mathematical Software* **18**, 360–373.

See Also

[gamma](#), [beta](#); my own [logQab_asy\(\)](#).

Examples

```
Qab <- algdiv(2:3, 8:14)
cbind(a = 2:3, b = 8:14, Qab) # recycling with a warning

## algdiv() and my logQab_asy() give *very* similar results for largish b:
all.equal( - algdiv(3, 100),
           logQab_asy(3, 100), tol=0) # 1.283e-16 !!
(lQab <- logQab_asy(3, 1e10))
## relative error
1 + lQab/ algdiv(3, 1e10) # 0 (64b F 30 Linux; 2019-08-15)
```

b_chi

Compute $E[\chi_{\nu}]/\sqrt{\nu}$ useful for t- and chi-Distributions

Description

$$b_{\chi}(\nu) := E[\chi(\nu)]/\sqrt{\nu} = \frac{\sqrt{2/\nu}\Gamma((\nu+1)/2)}{\Gamma(\nu/2)},$$

where $\chi(\nu)$ denotes a chi-distributed random variable, i.e., the square of a chi-squared variable, and $\Gamma(z)$ is the Gamma function, [gamma\(\)](#) in R.

This is a relatively important auxiliary function when computing with non-central t distribution functions and approximations, specifically see Johnson et al.(1994), p.520, after (31.26a), e.g., our [pntJW39\(\)](#).

Its logarithm,

$$lb_{\chi}(\nu) := \log\left(\frac{\sqrt{2/\nu}\Gamma((\nu+1)/2)}{\Gamma(\nu/2)}\right),$$

is even easier to compute via [lgamma](#) and [log](#), and I have used Maple to derive an asymptotic expansion in $\frac{1}{\nu}$ as well.

Note that $lb_{\chi}(\nu)$ also appears in the formula for the t-density ([dt](#)) and distribution (tail) functions.

Usage

```
b_chi      (nu, one.minus = FALSE, c1 = 341, c2 = 1000)
b_chiAsymp(nu, order = 2, one.minus = FALSE)
#lb_chi    (nu, ....) # not yet
lb_chiAsymp(nu, order)

c_dt(nu)    # warning("FIXME: current c_dt() is poor -- base it on lb_chi(nu) !")
c_dtAsymp(nu) # deprecated in favour of lb_chi(nu)
c_pt(nu)    # warning("use better c_dt()") %--> FIXME deprecate even stronger ?
```

Arguments

nu	non-negative numeric vector of degrees of freedom.
one.minus	logical indicating if $1 - b()$ should be returned instead of $b()$.
c1, c2	boundaries for different approximation intervals used: for $0 < \text{nu} \leq c1$, internal b1() is used, for $c1 < \text{nu} \leq c2$, internal b2() is used, and for $c2 < \text{nu}$, the b_chiAsymp() function is used, (and you can use that explicitly, also for smaller nu).
order	<p>FIXME: c1 and c2 were defined when the only asymptotic expansion known to me was the order = 2 one. A future version of b_chi will <i>very likely</i> use b_chiAsymp(*,order) for higher orders, and the c1 and c2 arguments will change, possibly be abolished.</p> <p>the polynomial order in $\frac{1}{\nu}$ of the asymptotic expansion of $b_\chi(\nu)$ for $\nu \rightarrow \infty$. The default, order = 2 corresponds to the order you can get out of the Abramowitz and Stegun (6.1.47) formula. Higher order expansions were derived using Maple by Martin Maechler in 2002, see below, but implemented in b_chiAsymp() only in 2018.</p>

Details

One can see that b_chi() has the properties of a CDF of a continuous positive random variable: It grows monotonely from $b_\chi(0) = 0$ to (asymptotically) one. Specifically, for large nu, b_chi(nu) = b_chiAsymp(nu) and

$$1 - b_\chi(\nu) \sim \frac{1}{4\nu}.$$

More accurately, derived from Abramowitz and Stegun, 6.1.47 (p.257) for a= 1/2, b=0,

$$\Gamma(z + 1/2)/\Gamma(z) \sim \sqrt{z} * (1 - 1/(8z) + 1/(128z^2) + O(1/z^3)),$$

and applied for $b_\chi(\nu)$ with $z = \nu/2$, we get

$$b_\chi(\nu) \sim 1 - (1/(4\nu) * (1 - 1/(8\nu)) + O(\nu^{-3})),$$

which has been implemented in b_chiAsymp(*,order=2) in 1999.

Even more accurately, Martin Maechler, used Maple to derive an asymptotic expansion up to order 15, here reported up to order 5, namely with $r := \frac{1}{4\nu}$,

$$b_\chi(\nu) = c_\chi(r) = 1 - r + \frac{1}{2}r^2 + \frac{5}{2}r^3 - \frac{21}{8}r^4 - \frac{399}{8}r^5 + O(r^6).$$

Value

a numeric vector of the same length as nu.

Author(s)

Martin Maechler

References

Johnson, Kotz, Balakrishnan (1995) *Continuous Univariate Distributions*, Vol 2, 2nd Edition; Wiley.

Formula on page 520, after (31.26a)

Abramowitz, M. and Stegun, I. A. (1972) *Handbook of Mathematical Functions*. New York: Dover. https://en.wikipedia.org/wiki/Abramowitz_and_Stegun provides links to the full text which is in public domain.

See Also

The t-distribution (base R) page [pt](#); our [pntJW39\(\)](#).

Examples

```
curve(b_chi, 0, 20); abline(h=0:1, v=0, lty=3)
r <- curve(b_chi, 1e-10, 1e5, log="x")
with(r, lines(x, b_chi(x, one.minus=TRUE), col = 2))

## Zoom in to c1-region
rc1 <- curve(b_chi, 340.5, 341.5, n=1001)# nothing to see
e <- 1e-3; curve(b_chi, 341-e, 341+e, n=1001) # nothing
e <- 1e-5; curve(b_chi, 341-e, 341+e, n=1001) # see noise, but no jump
e <- 1e-7; curve(b_chi, 341-e, 341+e, n=1001) # see float "granularity"+"jump"

## Zoom in to c2-region
rc2 <- curve(b_chi, 999.5, 1001.5, n=1001) # nothing visible
e <- 1e-3; curve(b_chi, 1000-e, 1000+e, n=1001) # clear small jump
c2 <- 1500
e <- 1e-3; curve(b_chi(x,c2=c2), c2-e, c2+e, n=1001)# still
## - - -
c2 <- 3000
e <- 1e-3; curve(b_chi(x,c2=c2), c2-e, c2+e, n=1001)# ok asympt clearly better!!
curve(b_chiAsymp, add=TRUE, col=adjustcolor("red", 1/3), lwd=3)
if(requireNamespace("Rmpfr")) {
  xm <- Rmpfr::seqMpfr(c2-e, c2+e, length.out=1000)

}

## - - -
c2 <- 4000
e <- 1e-3; curve(b_chi(x,c2=c2), c2-e, c2+e, n=1001)# ok asympt clearly better!!
curve(b_chiAsymp, add=TRUE, col=adjustcolor("red", 1/3), lwd=3)

grCol <- adjustcolor("forest green", 1/2)
curve(b_chi, 1/2, 1e11, log="x")
curve(b_chiAsymp, add = TRUE, col = grCol, lwd = 3)
## 1-b(nu) ~ 1/(4 nu) a power function <=> linear in log-log scale:
curve(b_chi(x, one.minus=TRUE), 1/2, 1e11, log="xy")
curve(b_chiAsymp(x, one.minus=TRUE), add = TRUE, col = grCol, lwd = 3)
```

dchisqApproxApproximations of the (Noncentral) Chi-Squared Density

Description

Compute the density function $f(x, *)$ of the (noncentral) chi-squared distribution.

Usage

```
dnchisqR      (x, df, ncp, log = FALSE,
                 eps = 5e-15, termSml = 1e-10, ncpLarge = 1000)
dnchisqBessel(x, df, ncp, log = FALSE)
dchisqAsym   (x, df, ncp, log = FALSE)
dnoncentchisq(x, df, ncp, kmax = floor(ncp/2 + 5 * (ncp/2)^0.5))
```

Arguments

<code>x</code>	non-negative numeric vector.
<code>df</code>	degrees of freedom (parameter), a positive number.
<code>ncp</code>	non-centrality parameter δ ;
<code>log</code>	logical indicating if the result is desired on the log scale.
<code>eps</code>	positive convergence tolerance for the series expansion: Terms are added while $\text{term} * q > (1-q)*\text{eps}$, where q is the term's multiplication factor.
<code>termSml</code>	positive tolerance: in the series expansion, terms are added to the sum as long as they are not smaller than <code>termSml * sum</code> even when convergence according to <code>eps</code> had occurred. This was not part of the original C code, but was added later for safeguarding against infinite loops, from PR#14105, e.g., for <code>dchisq(2000, 2, 1000)</code> .
<code>ncpLarge</code>	in the case where <code>mid</code> underflows to 0, when <code>log</code> is true, or <code>ncp >= ncpLarge</code> , use a central approximation. In theory, an optimal choice of <code>ncpLarge</code> would not be arbitrarily set at 1000 (hardwired in R's <code>dchisq()</code> here), but possibly also depend on <code>x</code> or <code>df</code> .
<code>kmax</code>	the number of terms in the sum for <code>dnoncentchisq()</code> .

Details

`dnchisqR()` is a pure R implementation of R's own C implementation in the sources, 'R/src/nmath/dnchisq.c', additionally exposing the three "tuning parameters" `eps`, `termSml`, and `ncpLarge`.

`dnchisqBessel()` implements Fisher(1928)'s exact closed form formula based on the Bessel function I_{nu} , i.e., R's `besselI()` function; specifically formula (29.4) in Johnson et al. (1995).

`dchisqAsym()` is the simple asymptotic approximation from Abramowitz and Stegun's formula 26.4.27, p. 942.

`dnoncentchisq()` uses the (typically defining) infinite series expansion directly, with truncation at `kmax`, and terms t_k which are products of a Poisson probability and a central chi-square density, i.e., terms `t.k := dpois(k, lambda = ncp/2) * dchisq(x, df = 2*k + df)` for $k = 0, 1, \dots, kmax$.

Value

numeric vector similar to `x`, containing the (logged if `log=TRUE`) values of the density $f(x, *)$.

Note

These functions are mostly of historical interest, notably as R's `dchisq()` was not always very accurate in the noncentral case, i.e., for $ncp > 0$.

Note

R's `dchisq()` is typically more uniformly accurate than the approximations nowadays, apart from `dnchisqR()` which should behave the same. There may occasionally exist small differences between `dnchisqR(x, *)` and `dchisq(x, *)` for the same parameters.

Author(s)

Martin Maechler, April 2008

References

Abramowitz, M. and Stegun, I. A. (1972) *Handbook of Mathematical Functions*. New York: Dover. https://en.wikipedia.org/wiki/Abramowitz_and_Stegun provides links to the full text which is in public domain.

Johnson, N.L., Kotz, S. and Balakrishnan, N. (1995) Continuous Univariate Distributions Vol~2, 2nd ed.; Wiley.

Chapter 29, Section 3 *Distribution*, (29.4), p. 436.

See Also

R's own `dchisq()`.

Examples

```
x <- sort(outer(c(1,2,5), 2^{(-4:5)}))
fRR <- dchisq (x, 10, 2)
f.R <- dnchisqR(x, 10, 2)
all.equal(fRR, f.R, tol = 0) # 64bit Lnx (F 30): 1.723897e-16
stopifnot(all.equal(fRR, f.R, tol = 4e-15))
```

Description

R transcriptions of the C code utility functions for `dgamma()` and similar “base” density functions by Catherine Loader.

Usage

```
dpois_raw(x, lambda, log)
bd0(x, np, verbose = getOption("verbose"))
stirlerr(n)
```

Arguments

<code>x, n</code>	a number (numeric).
<code>lambda, np</code>	a number (numeric); distribution parameters.
<code>log</code>	logical indicating if the log-density should be returned, otherwise the density at <code>x</code> .
<code>verbose</code>	logical indicating if some information about the computations are to be printed.

Value

a number

Author(s)

Martin Maechler

See Also

[dgamma](#), [dpois](#).

Examples

```
n <- seq(1, 50, by=1/4)
plot(n, sapply(n, stirlerr), type = "b", log="xy")
```

`dgamma.R`

Gamma Density Function Alternatives

Description

`dgamma.R()` is aimed to be an R level “clone” of R’s C level implementation [dgamma](#) (from package [stats](#)).

Usage

```
dgamma.R(x, shape, scale = 1, log)
```

Arguments

<code>x</code>	non-negative numeric vector.
<code>shape</code>	non-negative shape parameter of the Gamma distribution.
<code>scale</code>	positive scale parameter; note we do not see the need to have a <code>rate</code> parameter as the standard R function.
<code>log</code>	logical indicating if the result is desired on the log scale.

Value

numeric vector of the same length as `x` (which may have to be thought of recycled along `shape` and/or `scale`).

Author(s)

Martin Maechler

See Also

(As R's C code) this depends crucially on the “workhorse” function [dpois_raw\(\)](#).

Examples

```
## TODO: ... regular case .. use all.equal() ...

## PR#17577 - dgamma(x, shape)  for shape < 1 (=> +Inf at x=0) and very small x
stopifnot(exprs = {
  all.equal(dgamma.R(2^-1027, shape = .99, log=TRUE), 7.1127667376, tol=1e-10)
  all.equal(dgamma.R(2^-1031, shape = 1e-2, log=TRUE), 702.8889158, tol=1e-10)
  all.equal(dgamma.R(2^-1048, shape = 1e-7, log=TRUE), 710.30007699, tol=1e-10)
  all.equal(dgamma.R(2^-1048, shape = 1e-7, scale = 1e-315, log=TRUE),
            709.96858768, tol=1e-10)
})
## R's dgamma() gave all Inf in R <= 3.6.1 [and still there in 32-bit Windows !]
```

Description

`dntJKBf1` implements the summation formulas of Johnson, Kotz and Balakrishnan (1995), (31.15) on page 516 and (31.15') on p.519, the latter being typo-corrected for a missing factor $1/j!$.

`dntJKBf()` is [Vectorize](#)(`dntJKBf1`, c("x", "df", "ncp")), i.e., works vectorized in all three main arguments `x`, `df` and `ncp`.

The functions `.dntJKBch1()` and `.dntJKBch()` are only there for didactical reasons allowing to check that indeed formula (31.15') in the reference is missing a $j!$ factor in the denominator.

The `dntJKBf*`() functions are written to also work with arbitrary precise numbers of [class "mpfr"](#) (from package [Rmpfr](#)) as arguments.

Usage

```
dntJKBf1(x, df, ncp, log = FALSE, M = 1000)
dntJKBf (x, df, ncp, log = FALSE, M = 1000)

## The "checking" versions, only for proving correctness of formula:
.dntJKBch1(x, df, ncp, log = FALSE, M = 1000, check=FALSE, tol.check = 1e-7)
.dntJKBch (x, df, ncp, log = FALSE, M = 1000, check=FALSE, tol.check = 1e-7)
```

Arguments

x, df, ncp	see R's dt() ; note that each can be of class " mpfr ".
log	as in dt() , a logical indicating if $\log(f(x, *))$ should be returned instead of $f(x, *)$.
M	the number of terms to be used, a positive integer.
check	logical indicating if checks of the formula equalities should be done.
tol.check	tolerance to be used for all.equal() when check is true.

Details

How to choose M optimally has not been investigated yet.

Note that relatedly,

R's source code 'R/src/nmath/dnt.c' has claimed from 2003 till 2014 but **wrongly** that the non-central t density $f(x, *)$ is

```
f(x, df, ncp) =
df^(df/2) * exp(-.5*ncp^2) /
(sqrt(pi)*gamma(df/2)*(df+x^2)^((df+1)/2)) *
sum_{k=0}^Inf gamma((df + k + df)/2)*ncp^k / prod(1:k)*(2*x^2/(df+x^2))^(k/2) .
```

These functions (and this help page) prove that it was wrong.

Value

a number for [dntJKBf1\(\)](#) and [.dntJKBch1\(\)](#).

a numeric vector of the same length as the maximum of the lengths of x, df, ncp for [dntJKBf\(\)](#) and [.dntJKBch\(\)](#).

Author(s)

Martin Maechler

References

Johnson, N.L., Kotz, S. and Balakrishnan, N. (1995) Continuous Univariate Distributions Vol-2, 2nd ed.; Wiley.
Chapter 31, Section 5 *Distribution Function*, p.514 ff

See Also

R's [dt](#); (an improved version of) Viechtbauer's proposal: [dtWV](#).

Examples

```
tt <- seq(0, 10, len = 21)
ncp <- seq(0, 6, len = 31)
dt3R <- outer(tt, ncp, dt, df = 3)
dt3JKB <- outer(tt, ncp, dntJKBf, df = 3)
all.equal(dt3R, dt3JKB) # Lnx(64-b): 51 NA's in dt3R

x <- seq(-1,12, by=1/16)
fx <- dt(x, df=3, ncp=5)
```

```

re1 <- 1 - .dntJKBch(x, df=3, ncp=5) / fx ; summary(warnings()) # slow, with warnings
op <- options(warn = 2) # (=> warning == error, for now)
re2 <- 1 - dntJKBf (x, df=3, ncp=5) / fx # faster, no warnings
stopifnot(all.equal(re1[!is.na(re1)], re2[!is.na(re1)], tol=1e-6))
head( cbind(x, fx, re1, re2) , 20)
matplot(x, log10(abs(cbind(re1, re2))), type = "o", cex = 1/4)

## One of the numerical problems in "base R"'s non-central t-density:
options(warn = 0) # (factory def.)
x <- 2^seq(-12, 32, by=1/8) ; df <- 1/10
dtm <- cbind(dt(x, df=df, log=TRUE),
              dt(x, df=df, ncp=df/2, log=TRUE),
              dt(x, df=df, ncp=df, log=TRUE),
              dt(x, df=df, ncp=df*2, log=TRUE)) #.. quite a few warnings:
summary(warnings())
matplot(x, dtm, type="l", log = "x", xaxt="n",
        main = "dt(x, df=1/10, log=TRUE) central and noncentral")
sfsmisc::eaxis(1)
legend("right", legend=c("", paste0("ncp = df",c("/2","","*2"))),
       lty=1:4, col=1:4, bty="n")
# using MPFR high accuracy arithmetic (too slow for routine testing)
## no such kink here:
x. <- if(requireNamespace("Rmpfr")) Rmpfr::mpfr(x, 256) else x
system.time(dtJKB <- dntJKBf(x., df=df, ncp=df, log=TRUE)) # 7 sec if(Rmpfr)
lines(x, dtJKB, col=adjustcolor(3, 1/2), lwd=3)
options(op) # reset to prev.

## Relative Difference / Approximation errors :
plot(x, 1 - dtJKB / dtm[,3], type="l", log="x")
plot(x, 1 - dtJKB / dtm[,3], type="l", log="x", xaxt="n", ylim=c(-1,1)*1e-3); sfsmisc::eaxis(1)
plot(x, 1 - dtJKB / dtm[,3], type="l", log="x", xaxt="n", ylim=c(-1,1)*1e-7); sfsmisc::eaxis(1)
plot(x, abs(1 - dtJKB / dtm[,3]), type="l", log="xy", axes=FALSE, main =
      "dt(*, 1/10, 1/10, log=TRUE) relative approx. error",
      sub= paste("Copyright © 2019 Martin Mächler --- ", R.version.string))
for(j in 1:2) sfsmisc::eaxis(j)

```

Description

Compute the density function $f(x)$ of the t distribution with df degrees of freedom and non-centrality parameter ncp , according to Wolfgang Viechtbauer's proposal in 2002.

Usage

```
dtWV(x, df, ncp = 0, log = FALSE)
```

Arguments

x	numeric vector.
df	degrees of freedom (> 0 , maybe non-integer). $df = \text{Inf}$ is allowed.

ncp	non-centrality parameter δ ; If omitted, use the central t distribution.
log	logical; if TRUE, $\log(f(x))$ is returned instead of $f(x)$.

Details

The formula used is “asymptotic”: Resnikoff and Lieberman (1957), p.1 and p.25ff, proposed to use recursive polynomials for (*integer !*) degrees of freedom $f = 1, 2, \dots, 20$, and then, for $df = f > 20$, use the asymptotic approximation which Wolfgang Viechtbauer proposed as a first version of a non-central t density for R (when `dt()` did not yet have an ncp argument).

Value

numeric vector of density values, properly recycled in (x, df, ncp).

Author(s)

Wolfgang Viechtbauer (2002) post to R-help (<https://stat.ethz.ch/pipermail/r-help/2002-October/026044.html>), and Martin Maechler (log argument; tweaks, notably recycling).

References

Resnikoff, George J. and Lieberman, Gerald J. (1957) *Tables of the non-central t-distribution*; Technical report no. 32 (LIE ONR 32), April 1, 1957; Applied Math. and Stat. Lab., Stanford University.
<https://statistics.stanford.edu/research/tables-non-central-t-distribution-density-function-cum>

See Also

`dt`, R’s (C level) implementation of the (non-central) t density; `dntJKBF`, for Johnson et al.’s summation formula approximation.

Examples

```
tt <- seq(0, 10, len = 21)
ncp <- seq(0, 6, len = 31)
dt3R <- outer(tt, ncp, dt , df = 3)
dt3WV <- outer(tt, ncp, dtWV, df = 3)
all.equal(dt3R, dt3WV) # rel.err 0.00063
dt25R <- outer(tt, ncp, dt , df = 25)
dt25WV <- outer(tt, ncp, dtWV, df = 25)
all.equal(dt25R, dt25WV) # rel.err 1.1e-5

x <- -10:700
fx <- dt (x, df = 22, ncp =100)
lfx <- dt (x, df = 22, ncp =100, log=TRUE)
lfV <- dtWV(x, df = 22, ncp =100, log=TRUE)

head(lfx, 20) # shows that R's dt(*, log=TRUE) implementation is "quite suboptimal"

## graphics
opa <- par(no.readonly=TRUE)
par(mar=.1+c(5,4,4,3), mgp = c(2, .8,0))
plot(fx ~ x, type="l")
par(new=TRUE) ; cc <- c("red", adjustcolor("orange", 0.4))
plot(lfx ~ x, type = "o", pch=".", col=cc[1], cex=2, ann=FALSE, yaxt="n")
sfsmisc::eaxis(4, col=cc[1], col.axis=cc[1], small.args = list(col=cc[1]))
```

```

lines(x, lfv, col=cc[2], lwd=3)
dtt1 <- "      dt"; dtt2 <- "(x, df=22, ncp=100"; dttL <- paste0(dtt2, ", log=TRUE")
legend("right", c(paste0(dtt1,dtt2,")), paste0(c(dtt1,"dtWV"), dttL)),
      lty=1, lwd=c(1,1,3), col=c("black", cc), bty = "n")
par(opa) # reset

```

format01prec*Format Numbers in [0,1] with "Precise" Result***Description**

Format numbers in [0,1] with “precise” result, notably using “1-..” if needed.

Usage

```
format01prec(x, digits = getOption("digits"), width = digits + 2,
             eps = 1e-06, ...,
             FUN = function(x, ...) formatC(x, flag = "-", ...))
```

Arguments

<code>x</code>	numbers in [0,1]; (still works if not)
<code>digits</code>	number of digits to use; is used as <code>FUN(*, digits = digits)</code> or <code>FUN(*, digits = digits -5)</code> depending on <code>x</code> or <code>eps</code> .
<code>width</code>	desired width (of strings in characters), is used as <code>FUN(*, width = width)</code> or <code>FUN(*, width = width -2)</code> depending on <code>x</code> or <code>eps</code> .
<code>eps</code>	small positive number: Use '1-' for those <code>x</code> which are in $(1 - \text{eps}, 1]$. The author has claimed in the last millennium that (the default) 1e-6 is <i>optimal</i> .
<code>...</code>	optional further arguments passed to <code>FUN(x, digits, width, ...)</code> .
<code>FUN</code>	a function used for <code>format()</code> ing; must accept both a <code>digits</code> and <code>width</code> argument.

Value

a **character** vector of the same length as `x`.

Author(s)

Martin Maechler, 14 May 1997

See Also

`formatC`, `format.pval`.

Examples

```
## Show that format01prec() does reveal more precision :
cbind(format      (1 - 2^{-(16:24)}),
      format01prec(1 - 2^{-(16:24)}))

## a bit more variety
e <- c(2^seq(-24,0, by=2), 10^{-(7:1)})
ee <- sort(unique(c(e, 1-e)))
noquote(ff <- format01prec(ee))
data.frame(ee, format01prec = ff)
```

lbeta

(Log) Beta Approximations

Description

Compute $\log(\text{beta}(a,b))$ in a simple (fast) or asymptotic way.

Usage

```
lbetaM  (a, b, k.max = 5, give.all = FALSE)
lbeta_asy(a, b, k.max = 5, give.all = FALSE)
lbetaMM (a, b, cutAsy = 1e-2, verbose = FALSE)

betaI(a, n)
lbetaI(a, n)

logQab_asy(a, b, k.max = 5, give.all = FALSE)
Qab_terms(a, k)
```

Arguments

a, b, n	the Beta parameters, see <code>beta</code> ; n must be a positive integer and “small”.
k.max	..
give.all	<code>logical</code> ..
cutAsy	cutoff value from where to switch to asymptotic formula.
verbose	logical (or integer) indicating if and how much monitoring information should be printed to the console.
k	the number of terms in the series expansion of <code>Qab_terms()</code> , currently must be in {0, 1, .., 5}.

Details

All `lbeta*` functions compute $\log(\text{beta}(a,b))$.

We use $Q_{ab} = Q_{ab}(a,b)$ for

$$Q_{a,b} := \frac{\Gamma(a+b)}{\Gamma(b)},$$

which is numerically challenging when b becomes large compared to a , or $a \ll b$.

With the beta function

$$B(a, b) = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)} = \frac{\Gamma(a)}{Qab},$$

and hence

$$\log B(a, b) = \log \Gamma(a) + \log \Gamma(b) - \log \Gamma(a+b) = \log \Gamma(a) - \log Qab,$$

or in R, `lBeta(a,b) := lgamma(a) -logQab(a,b)`.

Indeed, typically everything has to be computed in log scale, as both $\Gamma(b)$ and $\Gamma(a+b)$ would overflow numerically for large b . Consequently, we use `logQab*`(\cdot), and for the large b case `logQab_asy`(\cdot) specifically,

$$\text{logQab}(a, b) := \log(Qab(a, b)).$$

Note this is related to trying to get asymptotic formula for Γ ratios, notably formula (6.1.47) in Abramowitz and Stegun.

Note how this is related to computing `qbeta()` in boundary cases, and see `algdiv()` ‘Details’ about this.

We also have a vignette about this, but really the problem has been addressed pragmatically by the authors of TOMS 708, see the ‘References’ in `pbeta`, by their routine `algdiv()` which also is available in our package **DPQ**.

Value

a fast or simple (approximate) computation of `lbeta(a,b)`.

Author(s)

Martin Maechler

References

Abramowitz, M. and Stegun, I. A. (1972) *Handbook of Mathematical Functions*. New York: Dover. https://en.wikipedia.org/wiki/Abramowitz_and_Stegun provides links to the full text which is in public domain.

Formula (6.1.47), p.257

See Also

R’s `beta` function; `algdiv()`.

Examples

```
## TODO
```

lgamma1p	<i>Accurate log(gamma(a+1))</i>
----------	---------------------------------

Description

Compute

$$l\Gamma_1(a) := \log \Gamma(a + 1) = \log(a \cdot \Gamma(a)) = \log a + \log \Gamma(a),$$

which is “in principle” the same as `log(gamma(a+1))` or `lgamma(a+1)`, accurately also for (very) small a ($0 < a < 0.5$).

Usage

```
lgamma1p (a, tol_logcf = 1e-14)
lgamma1p.(a, cutoff.a = 1e-6, k = 3)
lgamma1p_series(x, k)
```

Arguments

<code>a, x</code>	a numeric vector.
<code>tol_logcf</code>	for <code>lgamma1p()</code> : a non-negative number ...
<code>cutoff.a</code>	for <code>lgamma1p.()</code> : a positive number indicating the cutoff to switch from ...
<code>k</code>	an integer, the number of terms in the series expansion used internally.

Details

`lgamma1p()` is an R translation of the function (in Fortran) in Didonato and Morris (1992) which uses a 40-degree polynomial approximation.

`lgamma1p_series(x,k)` is Taylor series approximation of order k , (derived via Maple), which is $-\gamma x + \pi^2 x^2/12 + O(x^3)$, where γ is Euler’s constant 0.5772156649....

Value

a numeric vector with the same attributes as `a`.

Author(s)

Morten Welinder (C code of Jan 2005, see R’s bug issue PR#7307) for `lgamma1p()`.

Martin Maechler, notably for `lgamma1p_series()` which works with package **Rmpfr** but otherwise may be *much* less accurate than Morten’s 40 term series!

References

Didonato, A. and Morris, A., Jr, (1992) Algorithm 708: Significant digit computation of the incomplete beta function ratios. *ACM Transactions on Mathematical Software*, **18**, 360–373; see also `pbeta`.

See Also

`log1pmx`, `log1p`, `pbeta`.

Examples

```
curve(-log(x*gamma(x)), 1e-30, .8, log="xy", col="gray50", lwd = 3,
      axes = FALSE, ylim = c(1e-30,1))
sfsmisc::axis(1); sfsmisc::axis(2)
at <- 10^(1-4*(0:8))
abline(h = at, v = at, col = "lightgray", lty = "dotted")
curve(-lgamma( 1+x), add=TRUE, col="red2", lwd=1/2)# underflows even earlier
curve(-lgamma1p (x), add=TRUE, col="blue")
curve(-lgamma1p.(x), add=TRUE, col=adjustcolor("forest green",1/4),
      lwd = 5, lty = 2)
for(k in 1:7)
  curve(-lgamma1p_series(x, k=k), add=TRUE, col=paste0("gray",30+k*8), lty = 3)
```

log1mexp

Compute $f(a) = \log(1 - \exp(-a))$ Numerically Optimally

Description

Compute $f(a) = \log(1 - \exp(-a))$ quickly and numerically accurately.

Usage

`log1mexp(x)`

Arguments

`x` numeric vector of positive values.

Author(s)

Martin Maechler

References

Martin Mächler (2012). Accurately Computing $\log(1 - \exp(-|a|))$; <https://CRAN.R-project.org/package=Rmpfr/vignettes/log1mexp-note.pdf>.

See Also

The `log1mexp()` function in CRAN package **copula**, and the corresponding vignette (in the ‘References’).

<code>log1pmx</code>	<i>Accurate log(1+x) - x</i>
----------------------	------------------------------

Description

Compute

$$\log(1 + x) - x$$

accurately also for small x , i.e., $|x| \ll 1$.

Usage

```
log1pmx(x, tol_logcf = 1e-14)
```

Arguments

- | | |
|------------------------|--|
| <code>x</code> | numeric vector with values $x > -1$. |
| <code>tol_logcf</code> | a non-negative number indicating the tolerance (maximal relative error) for the auxiliary <code>logcf()</code> function. |

Details

In order to provide full accuracy, the computations happens differently in three regions for x ,

$$m_l = -0.79149064$$

is the first cutpoint,

- $x < m_l$ or $x > 1$: use `log1pmx(x) := log1p(x) - x`,
 - $|x| < 0.01$: use $t(((2/9 * y + 2/7)y + 2/5)y + 2/3)y - x$,
 - $x \in [m_l, 1]$, and $|x| \geq 0.01$: use $t(2ylogcf(y, 3, 2) - x)$,
- where $t := \frac{x}{2+x}$, and $y := t^2$.

Note that the formulas based on t are based on the (fast converging) formula

$$\log(1 + x) = 2 \left(r + \frac{r^3}{3} + \frac{r^5}{5} + \dots \right),$$

where $r := x/(x + 2)$, see the reference.

Value

a numeric vector (with the same attributes as `x`).

Author(s)

A translation of Morten Welinder's C code of Jan 2005, see R's bug issue [PR#7307](#).

References

- Abramowitz, M. and Stegun, I. A. (1972) *Handbook of Mathematical Functions*. New York: Dover.
https://en.wikipedia.org/wiki/Abramowitz_and_Stegun provides links to the full text which is in public domain.
 Formula (4.1.29), p.68.

See Also

[logcf](#), the auxiliary function, [lgamma1p](#) which calls [log1pmx](#), [log1p](#)

Examples

```
l1x <- curve(log1pmx, -.9999, 7, n=1001)
abline(h=0, v=-1:0, lty=3)
l1xz <- curve(log1pmx, -.1, .1, n=1001); abline(h=0, v=0, lty=3)
l1xz2 <- curve(log1pmx, -.01, .01, n=1001); abline(h=0, v=0, lty=3)
l1xz3 <- curve(-log1pmx(x), -.002, .002, n=2001, log="y", yaxt="n")
sfsmisc::eaxis(2); abline(v=0, lty=3)
```

logcf*Continued Fraction Approximation of Log-Related Series***Description**

Compute a continued fraction approximation to the series (infinite sum)

$$\sum_{k=0}^{\infty} \frac{x^k}{i + k \cdot d} = \frac{1}{i} + \frac{x}{i + d} + \frac{x^2}{i + 2 * d} + \frac{x^3}{i + 3 * d} + \dots$$

Needed as auxiliary function in [log1pmx\(\)](#) and [lgamma1p\(\)](#).

Usage

```
logcf(x, i, d, eps, maxit = 10000)
```

Arguments

x	numeric vector
i	positive numeric
d	non-negative numeric
eps	positive number, the convergence tolerance.
maxit	a positive integer, the maximal number of iterations or terms in the truncated series used.

Value

a numeric vector with the same attributes as x.

Note

Rescaling is done by (namespace hidden) “global” scalefactor

Author(s)

Martin Maechler, based on

See Also

[lgamma1p](#), [log1pmx](#), and [pbeta](#), whose principal algorithm has evolved from TOMS 708.

Examples

```
132 <- curve(logcf(x, 3,2, eps=1e-7), -3, 1)
abline(h=0,v=1, lty=3, col="gray50")
plot(y~x, 132, log="y", type = "o", main = "logcf(*, 3,2) in log-scale")
```

`logspace.add`

Logspace Arithmetix – Addition and Subtraction

Description

Compute the log(arithm) of a sum (or difference) from the log of terms without causing overflows and without throwing away large handfuls of accuracy.

`logspace.add(lx, ly):=`

$$\log(\exp(lx) + \exp(ly))$$

`logspace.sub(lx, ly):=`

$$\log(\exp(lx) - \exp(ly))$$

Usage

```
logspace.add(lx, ly)
logspace.sub(lx, ly)
```

Arguments

<code>lx, ly</code>	numeric vectors, typically of the same length , but will be recycled to common length as with other R arithmetic.
---------------------	---

Value

a [numeric](#) vector of the same length as `x+y`.

Note

This is really from R's C source code for [pgamma\(\)](#), i.e., '<R>/src/nmath/pgamma.c'

The function definitions are very simple, `logspace.sub()` using [log1mexp\(\)](#).

Author(s)

Morten Welinder (for R's `pgamma()`); Martin Maechler

See Also

[lsum](#), [lssum](#); then [pgamma\(\)](#)

Examples

```
set.seed(12)
ly <- rnorm(100, sd= 50)
lx <- ly + abs(rnorm(100, sd=100)) # lx - ly must be positive for *.sub()
stopifnot(exprs = {
  all.equal(logspace.add(lx,ly),
            log(exp(lx) + exp(ly)), tol=1e-14)
  all.equal(logspace.sub(lx,ly),
            log(exp(lx) - exp(ly)), tol=1e-14)
})
}
```

lssum

Compute Logarithm of a Sum with Signed Large Summands

Description

Properly compute $\log(x_1 + \dots + x_n)$ for given log absolute values `lxabs = log(|x1|), .., log(|xn|)` and corresponding signs `signs = sign(x1), .., sign(xn)`. Here, x_i is of arbitrary sign.

Notably this works in many cases where the direct sum would have summands that had overflowed to `+Inf` or underflowed to `-Inf`.

This is a (simpler, vector-only) version of `copula:::lssum()` (CRAN package **copula**).

Note that the *precision* is often not the problem for the direct summation, as R's `sum()` internally uses "long double" precision on most platforms.

Usage

```
lssum(lxabs, signs, l.off = max(lxabs), strict = TRUE)
```

Arguments

<code>lxabs</code>	n-vector of values $\log(x_1), \dots, \log(x_n)$.
<code>signs</code>	corresponding signs $sign(x_1), \dots, sign(x_n)$.
<code>l.off</code>	the offset to subtract and re-add; ideally in the order of <code>max(.)</code> .
<code>strict</code>	<code>logical</code> indicating if the function should stop on some negative sums.

Value

$\log(x_1 + \dots + x_n) == \log(sum(x)) = \log(sum(sign(x)*|x|)) == \log(sum(sign(x)*exp(\log(|x|)))) == \log(exp(\log(sign(x)*|x|)))$

Author(s)

Marius Hofert and Martin Maechler (for package **copula**).

See Also

`lsum()` which computes an exponential sum in log scale with *out* signs.

Examples

```
rSamp <- function(n, lmean, lsd = 1/4, roundN = 16) {
  lax <- sort((1+1e-14*rnorm(n))*round(roundN*rnorm(n, m = lmean, sd = lsd))/roundN)
  sx <- rep_len(c(-1,1), n)
  list(lax=lax, sx=sx, x = sx*exp(lax))
}

set.seed(101)
L1 <- rSamp(1000, lmean = 700) # here, lssum() is not needed (no under-/overflow)
summary(as.data.frame(L1))
ax <- exp(lax <- L1$lax)
hist(lax); rug(lax)
hist(ax); rug(ax)
sx <- L1$sx
table(sx)
(lsSimple <- log(sum(L1$x)))           # 700.0373
(lsS <- lssum(lxabs = lax, signs = sx))# ditto
lsS - lsSimple # even exactly zero (in 64b Fedora 30 Linux which has nice 'long double')
stopifnot(all.equal(700.037327351478, lsS, tol=1e-14), all.equal(lsS, lsSimple))

L2 <- within(L1, { lax <- lax + 10; x <- sx*exp(lax) }) ; summary(L2$x) # some -Inf, +Inf
(lsSimpl2 <- log(sum(L2$x)))           # NaN
(lsS2 <- lssum(lxabs = L2$lax, signs = L2$sx)) # 710.0373
stopifnot(all.equal(lsS2, lsS + 10, tol = 1e-14))
```

lsum

Properly Compute the Logarithm of a Sum

Description

Simple vector version of copula:::lsum() (CRAN package **copula**).

Properly compute $\log(x_1 + \dots + x_n)$. for given $\log(x_1), \dots, \log(x_n)$. Here, $x_i > 0$ for all i .

Usage

```
lsum(lx, l.off = max(lx))
```

Arguments

lx	n-vector of values $\log(x_1), \dots, \log(x_n)$.
l.off	the offset to subtract and re-add; ideally in the order of the maximum of each column.

Value

$$\log(x_1 + \dots + x_n) = \log(\sum(x)) = \log(\sum(\exp(\log(x)))) == \log(\exp(\log(\max(x))) * \sum(\exp(\log(x) - \log(\max(x))))$$

Author(s)

Originally, via paired programming: Marius Hofert and Martin Maechler.

See Also

[lsum\(\)](#) which computes a sum in log scale with specified (typically alternating) signs.

Examples

```
## The "naive" version :
lsum0 <- function(lx) log(sum(exp(lx)))

lx1 <- 10*(-80:70) # is easy
lx2 <- 600:750 # lsum0() not ok [could work with rescaling]
lx3 <- -(750:900) # lsum0() = -Inf - not good enough
m3 <- cbind(lx1, lx2, lx3)
lx6 <- lx5 <- lx4 <- lx3
lx4[149:151] <- -Inf ## = log(0)
lx5[150] <- Inf
lx6[1] <- NA_real_
m6 <- cbind(m3, lx4, lx5, lx6)
stopifnot(exprs = {
  all.equal(lsum(lx1), lsum0(lx1))
  all.equal((ls1 <- lsum(lx1)), 700.000045400960403, tol=8e-16)
  all.equal((ls2 <- lsum(lx2)), 750.458675145387133, tol=8e-16)
  all.equal((ls3 <- lsum(lx3)), -749.541324854612867, tol=8e-16)
  ## identical: matrix-version <==> vector versions
  identical(lsum(lx4), ls3)
  identical(lsum(lx4), lsum(head(lx4, -3))) # the last three were -Inf
  identical(lsum(lx5), Inf)
  identical(lsum(lx6), lx6[1])
  identical((lm3 <- apply(m3, 2, lsum)), c(lx1=ls1, lx2=ls2, lx3=ls3))
  identical(apply(m6, 2, lsum), c(lm3, lx4=ls3, lx5=Inf, lx6=lx6[1]))
})
}
```

Description

Given the function `G()` and its derivative `g()`, `newton()` uses the Newton method, starting at `x0`, to find a point `xp` at which `G` is zero. `G()` and `g()` may each depend on the same parameter (vector) `z`.

Convergence typically happens when the stepsize becomes smaller than `eps`.

`keepAll = TRUE` to also get the vectors of consecutive values of `x` and `G(x, z)`;

Usage

```
newton(x0, G, g, z,
       xMin = -Inf, xMax = Inf, warnRng = TRUE,
       dxMax = 1000, eps = 0.0001, maxiter = 1000L,
       warnIter = missing(maxiter) || maxiter >= 10L,
       keepAll = NA)
```

Arguments

<code>x0</code>	numeric start value.
<code>G, g</code>	must be functions , mathematically of their first argument, but they can accept parameters; <code>g()</code> must be the derivative of <code>G</code> .
<code>z</code>	parameter vector for <code>G()</code> and <code>g()</code> , to be kept fixed.
<code>xMin, xMax</code>	numbers defining the allowed range for <code>x</code> during the iterations; e.g., useful to set to 0 and 1 during quantile search.
<code>warnRng</code>	logical specifying if a warning should be signalled when start value <code>x0</code> is outside <code>[xMin, xMax]</code> and hence will be changed to one of the boundary values.
<code>dxMax</code>	maximal step size in <i>x</i> -space. (The default 1000 is quite arbitrary, do set a good maximal step size yourself!)
<code>eps</code>	positive number, the <i>absolute</i> convergence tolerance.
<code>maxiter</code>	positive integer, specifying the maximal number of Newton iterations.
<code>warnIter</code>	logical specifying if a warning should be signalled when the algorithm has not converged in <code>maxiter</code> iterations.
<code>keepAll</code>	logical specifying if the full sequence of <code>x</code> - and <code>G(x,*)</code> values should be kept and returned:
	NA, the default: <code>newton</code> returns a small list of final “data”, with 4 components <code>x = x*</code> , <code>G = G(x*, z)</code> , <code>it</code> , and <code>converged</code> .
	TRUE: returns an extended list , in addition containing the vectors <code>x.vec</code> and <code>G.vec</code> .
	FALSE: returns only the <code>x*</code> value.

Details

Because of the quadratic convergence at the end of the Newton algorithm, often x^* satisfies approximately $|G(x^*, z)| < \text{eps}^2$.

`newton()` can be used to compute the quantile function of a distribution, if you have a good starting value, and provide the cumulative probability and density functions as R functions `G` and `g` respectively.

Value

The result always contains the final `x`-value `x*`, and typically some information about convergence, depending on the value of `keepAll`, see above:

<code>x</code>	the optimal x^* value (a number).
<code>G</code>	the function value $G(x^*, z)$, typically very close to zero.
<code>it</code>	the integer number of iterations used.
<code>convergence</code>	logical indicating if the Newton algorithm converged within <code>maxiter</code> iterations.
<code>x.vec</code>	the full vector of <code>x</code> values, $\{x_0, \dots, x^*\}$.
<code>G.vec</code>	the vector of function values (typically tending to zero), i.e., $G(x.\text{vec}, .)$ (even when $G(x, .)$ would not vectorize).

Author(s)

Martin Maechler, ca. 2004

References

Newton's Method on Wikipedia, https://en.wikipedia.org/wiki/Newton%27s_method.

See Also

`uniroot()` is much more sophisticated, works without derivatives and is generally faster than `newton()`.
`newton(.)` is currently crucially used (only) in our function `qchisqN()`.

Examples

```
## The most simple non-trivial case : Computing SQRT(a)
G <- function(x, a) x^2 - a
g <- function(x, a) 2*x

newton(1, G, g, z = 4) # z = a -- converges immediately
newton(1, G, g, z = 400) # bad start, needs longer to converge

## More interesting, and related to non-central (chisq, e.t.) computations:
## When is x * log(x) < B, i.e., the inverse function of G = x*log(x) :
xlx <- function(x, B) x*log(x) - B
dxlx <- function(x, B) log(x) + 1

Nxlx <- function(B) newton(B, G=xlx, g=dxlx, z=B, maxiter=Inf)$x
N1   <- function(B) newton(B, G=xlx, g=dxlx, z=B, maxiter = 1)$x
N2   <- function(B) newton(B, G=xlx, g=dxlx, z=B, maxiter = 2)$x

Bs <- c(outer(c(1,2,5), 10^(0:4)))
plot (Bs, vapply(Bs, Nxlx, pi), type = "l", log ="xy")
lines(Bs, vapply(Bs, N1 , pi), col = 2, lwd = 2, lty = 2)
lines(Bs, vapply(Bs, N2 , pi), col = 3, lwd = 3, lty = 3)

BL <- c(outer(c(1,2,5), 10^(0:6)))
plot (BL, vapply(BL, Nxlx, pi), type = "l", log ="xy")
lines(BL, BL, col="green2", lty=3)
lines(BL, vapply(BL, N1 , pi), col = 2, lwd = 2, lty = 2)
lines(BL, vapply(BL, N2 , pi), col = 3, lwd = 3, lty = 3)
## Better starting value from an approximate 1 step Newton:
iL1 <- function(B) 2*B / (log(B) + 1)
lines(BL, iL1(BL), lty=4, col="gray20") ## really better ==> use it as start

Nxlx <- function(B) newton(iL1(B), G=xlx, g=dxlx, z=B, maxiter=Inf)$x
N1   <- function(B) newton(iL1(B), G=xlx, g=dxlx, z=B, maxiter = 1)$x
N2   <- function(B) newton(iL1(B), G=xlx, g=dxlx, z=B, maxiter = 2)$x

plot (BL, vapply(BL, Nxlx, pi), type = "o", log ="xy")
lines(BL, iL1(BL), lty=4, col="gray20")
lines(BL, vapply(BL, N1 , pi), type = "o", col = 2, lwd = 2, lty = 2)
lines(BL, vapply(BL, N2 , pi), type = "o", col = 3, lwd = 2, lty = 3)
## Manual 2-step Newton
iL2 <- function(B) { lB <- log(B) ; B*(lB+1) / (lB * (lB - log(lB) + 1)) }
lines(BL, iL2(BL), col = adjustcolor("sky blue", 0.6), lwd=6)
##=> iL2() is very close to true curve
## relative error:
iLtrue <- vapply(BL, Nxlx, pi)
cbind(BL, iLtrue, iL2=iL2(BL), relErL2 = 1-iL2(BL)/iLtrue)
```

```

## absolute error (in log-log scale; always positive!):
plot(BL, iL2(BL) - iLtrue, type = "o", log="xy", axes=FALSE)
if(requireNamespace("sfsmisc")) {
  sfsmisc::eaxis(1)
  sfsmisc::eaxis(2, sub10=2)
} else {
  cat("no 'sfsmisc' package; maybe  install.packages(\"sfsmisc\")  ?\n")
  axis(1); axis(2)
}
## 1 step from iL2()  seems quite good:
B. <- BL[-1] # starts at 2
NL2 <- lapply(B., function(B) newton(iL2(B), G=xlx, g=dxlx, z=B, maxiter=1))
str(NL2)
iL3 <- sapply(NL2, `[[[`, "x")
cbind(B., iLtrue[-1], iL2=iL2(B.), iL3, relE.3 = 1- iL3/iLtrue[-1])
x. <- iL2(B.)
all.equal(iL3, x. - xlx(x., B.) / dxlx(x.)) ## 7.471802e-8
## Algebraic simplification of one newton step :
all.equal((x.+B.)/(log(x.)+1), x. - xlx(x., B.) / dxlx(x.), tol = 4e-16)
iN1 <- function(x, B) (x+B) / (log(x) + 1)
B <- 12345
iN1(iN1(iN1(B, B), B), B)
Nxlx(B)

```

Description

The **DPQ** package provides some numeric constants used in some of its distribution computations. `all_mpfr()` and `any_mpfr()` return `TRUE` iff all (or ‘any’, respectively) of their arguments inherit from class “`mpfr`” (from package **Rmpfr**).

`logr(x,a)` computes `log(x / (x + a))` in a numerically stable way.

Usage

```

## Numeric Constants : % mostly in  .../R/beta-fns.R
M_LN2      # = log(2) = 0.693...
M_cutoff   # := If |x| > |k| * M_cutoff, then  log[ exp(-x) * k^x ]  =~= -x
            # = 3196577161300663808 ~= 3.2e+18
M_minExp   # = log(2) * .Machine$double.min.exp # ~= -708.396..
G_half     # = sqrt(pi) = Gamma( 1/2 )

## Functions :
all_mpfr(...)
any_mpfr(...)
logr(x, a)  # == log(x / (x + a)) -- but numerically smart; x >= 0, a > -x
okLongDouble(lambda = 999, verbose = 0L, tol = 1e-15)

```

Arguments

...	numeric or "mpfr" numeric vectors.
x, a	number-like, not negative, now may be vectors of <code>length(.) > 1</code> .
lambda	a number, typically in the order of 500–10'000.
verbose	a non-negative integer, if not zero, <code>okLongDouble()</code> prints the intermediate long double computations' results.
tol	numerical tolerance used to determine the accuracy required for near equality in <code>okLongDouble()</code> .

Details

`all_mpfr()`,
`all_mpfr()` : test if `all` or `any` of their arguments or of class "mpfr" (from package **Rmpfr**). The arguments are evaluated only until the result is determined, see the example.
`logr()` computes $\log(x/(x + a))$ in a numerically stable way.

Value

The numeric constant in the first case; a numeric (or "mpfr") vector of appropriate size in the 2nd case.
`okLongDouble()` returns a `logical`, `TRUE` iff the long double arithmetic with `expl()` and `log1()` seems to work accurately

Author(s)

Martin Maechler

See Also

[.Machine](#)

Examples

```
(Ms <- ls("package:DPQ", pattern = "^M"))
lapply(Ms, function(nm) { cat(nm, ": "); print(get(nm)) }) -> .tmp

logr(1:3, a=1e-10)

okLongDouble() # typically TRUE, but not e.g. in a valgrinded R-devel of Oct.2019
## Here is typically the "boundary":
okLongDouble(11355, verbose=TRUE) # typically TRUE (also for lambda <= 11355)
okLongDouble(11356, verbose=TRUE) # typically FALSE (also for lambda >= 11356)
```

pbetaRv1

Pure R Implementation of Old pbeta()

Description

`pbetaRv1()` is an implementation of the original (“version 1” `pbeta()`) function in R (versions \leq 2.2.x), before we started using TOMS 708 `bratio()` instead, see that help page also for references.

`pbetaRv1()` is basically a manual translation from C to R of the underlying `pbeta_raw()` C function, see in R’s source tree at <https://svn.r-project.org/R/branches/R-2-2-patches/src/nmath/pbeta.c>

For consistency within R, we are using R’s argument names (`q, shape1, shape2`) instead of C code’s (`x, pin, qin`).

It is only for the *central* beta distribution.

Usage

```
pbetaRv1(q, shape1, shape2, lower.tail = TRUE,
          eps = 0.5 * .Machine$double.eps,
          sml = .Machine$double.xmin,
          verbose = 0)
```

Arguments

<code>q, shape1, shape2</code>	non-negative numbers, <code>q</code> in $[0, 1]$, see <code>pbeta</code> .
<code>lower.tail</code>	indicating if $F(q; *)$ should be returned or the upper tail probability $1 - F(q)$.
<code>eps</code>	the tolerance used to determine convergence. <code>eps</code> has been hard coded in C code to $0.5 * .Machine$double.eps$ which is equal to 2^{-53} or $1.110223e-16$.
<code>sml</code>	the smallest positive number on the typical platform. The default <code>.Machine\$double.xmin</code> is hard coded in the C code (as <code>DBL_MIN</code>), and this is equal to 2^{-1022} or $2.225074e-308$ on all current platforms.
<code>verbose</code>	integer indicating the amount of verbosity of diagnostic output, 0 means no output, 1 more, etc.

Value

a number.

Note

The C code contains

This routine is a translation into C of a Fortran subroutine by W. Fullerton of Los Alamos Scientific Laboratory.

Author(s)

Martin Maechler

References

(From the C code:)

Nancy E. Bosten and E.L. Battiste (1974). Remark on Algorithm 179 (S14): Incomplete Beta Ratio. *Communications of the ACM*, **17**(3), 156–7.

See Also

[pbeta](#).

Examples

```
all.equal(pbtaRv1(1/4, 2, 3),
          pbeta  (1/4, 2, 3))
set.seed(101)

N <- 1000
x <- sample.int(7, N, replace=TRUE) / 8
a <- rlnorm(N)
b <- 5*rlnorm(N)
pbt <- pbeta(x, a, b)
for(i in 1:N) {
  stopifnot(all.equal(pbtaRv1(x[i], a[i], b[i]), pbt[i]))
  cat(".", if(i %% 20 == 0) paste0(i, "\n"))
}
```

Description

Plot two noncentral (chi-squared or t or ..) distribution curves for visual comparison.

Usage

```
pl2curves(fun1, fun2, df, ncp, log = FALSE,
          from = 0, to = 2 * ncp, p.log = "", n = 2001,
          leg = TRUE, col2 = 2, lwd2 = 2, lty2 = 3, ...)
```

Arguments

fun1, fun2	<code>function()</code> s, both to be used via <code>curve()</code> , and called with the same 4 arguments, <code>(., df, ncp, log)</code> (the name of the first argument is not specified).
df, ncp, log	parameters to be passed and used in both functions, which hence typically are non-central chi-squared or t density, probability or quantile functions.
from, to	numbers determining the x-range, passed to <code>curve()</code> .
p.log	string, passed as <code>curve(..., log = log.p)</code> .
n	the number of evaluation points, passed to <code>curve()</code> .
leg	logical specifying if a <code>legend()</code> should be drawn.

```
col2, lwd2, lty2
color, line width and line type for the second curve. (The first curve uses defaults
for these graphical properties.)
...
further arguments passed to first curve(..) call.
```

Value

TODO: invisible return both curve() results, i.e., (x,y1, y2), possibly as data frame

Author(s)

Martin Maechler

See Also

[curve](#), ..

Examples

```
p.dnchiBessel <- function(df, ncp, log=FALSE, from=0, to = 2*ncp, p.log="", ...)
{
  pl2curves(dnchisqBessel, dchisq, df=df, ncp=ncp, log=log,
             from=from, to=to, p.log=p.log, ...)
}

## TODO the p.dnchiB() examples  >>>> ../../tests/chisq-nonc-ex.R <<<
```

Description

`pnbetaAppr2()` and its initial version `pnbetaAppr2v1()` provide the “approximation 2” of Chatamvelli and Shanmugam(1997) to the noncentral Beta probability distribution.

`pnbetaAS310()` is an R level interface to a C translation (and “Rification”) of the AS 310 Fortran implementation.

Usage

```
pnbetaAppr2(x, a, b, ncp = 0, lower.tail = TRUE, log.p = FALSE)

pnbetaAS310(x, a, b, ncp = 0, lower.tail = TRUE, log.p = FALSE,
             useAS226 = (ncp < 54.),
             errmax = 1e-6, itrmax = 100)
```

Arguments

<code>x</code>	numeric vector (of quantiles), typically from inside [0, 1].
<code>a, b</code>	the shape parameters of Beta, aka as <code>shape1</code> and <code>shape2</code> .
<code>ncp</code>	non-centrality parameter.
<code>log.p</code>	logical; if TRUE, probabilities p are given as log(p).
<code>lower.tail</code>	logical; if TRUE (default), probabilities are $P[X \leq x]$, otherwise, $P[X > x]$.
<code>useAS226</code>	<code>logical</code> specifying if AS 226 (with R84 and R95 amendments) should be used which is said to be sufficient for small ncp. The default <code>ncp < 54</code> had been hardwired in AS 310.
<code>errmax</code>	non-negative number determining convergence for AS 310.
<code>itrmax</code>	positive integer number, only if(<code>useAS226</code>) is passed to AS 226.

Value

a numeric vector of (log) probabilities of the same length as `x`.

Note

The authors in the reference compare AS 310 with Lam(1995), Frick(1990) and Lenth(1987) and state to be better than them. R's current (2019) noncentral beta implementation builds on these, too, with some amendments though; still, `pnbetaAS310()` may potentially be better, at least in certain corners of the 4-dimensional input space.

Author(s)

Martin Maechler; `pnbetaAppr2()` in Oct 2007.

References

- Chattamvelli, R., and Shanmugam, R. (1997) Algorithm AS 310: Computing the Non-Central Beta Distribution Function. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **46**(1), 146–156, for “approximation 2” notably p.154;
 doi: [10.1111/14679876.00055](https://doi.org/10.1111/14679876.00055).
- Lenth, R. V. (1987) Algorithm AS 226, ..., Frick, H. (1990)'s AS R84, ..., and Lam, M.L. (1995)'s AS R95 : See ‘References’ in R's `pbeta` page.

See Also

R's own `pbeta`.

Examples

```
## Same arguments as for Table 1 (p.151) of the reference
a <- 5*rep(1:3, each=3)
aargs <- cbind(a = a, b = a,
               ncp = rep(c(54, 140, 170), 3),
               x = 1e-4*c(8640, 9000, 9560, 8686, 9000, 9000, 8787, 9000, 9220))
aargs
pnA2 <- apply(aargs, 1, function(aa) do.call(pnbetaAppr2, as.list(aa)))
pnA310<- apply(aargs, 1, function(aa) do.call(pnbetaAS310, as.list(aa)))
aar2 <- aargs; dimnames(aar2)[[2]] <- c(paste0("shape", 1:2), "ncp", "q")
```

```

pnbR <- apply(aar2, 1, function(aa) do.call(pbta, as.list(aa)))
range(relD2 <- 1 - pnbA2 /pnbR)
range(relD310 <- 1 - pnA310/pnbR)
cbind(aargs, pnbA2, pnA310, pnbR,
      relD2 = signif(relD2, 3), relD310 = signif(relD310, 3)) # <-----> Table 1
stopifnot(abs(relD2) < 0.009) # max is 0.006286
stopifnot(abs(relD310) < 1e-5 ) # max is 6.3732e-6

## Arguments as for Table 2 (p.152) of the reference :
aarg2 <- cbind(a = c( 10, 10, 15, 20, 20, 30, 30),
                 b = c( 20, 10, 5, 10, 30, 50, 20, 40),
                 ncp=c(150,120, 80,110, 65,130, 80,130),
                 x = c(868,900,880,850,660,720,720,800)/1000)
pnbA2 <- apply(aarg2, 1, function(aa) do.call(pnbtAppr2, as.list(aa)))
pnA310<- apply(aarg2, 1, function(aa) do.call(pnbtAS310, as.list(aa)))
aar2 <- aarg2; dimnames(aar2)[[2]] <- c(paste0("shape", 1:2), "ncp", "q")
pnbR <- apply(aar2, 1, function(aa) do.call(pbta, as.list(aa)))
range(relD2 <- 1 - pnbA2 /pnbR)
range(relD310 <- 1 - pnA310/pnbR)
cbind(aarg2, pnbA2, pnA310, pnbR,
      relD2 = signif(relD2, 3), relD310 = signif(relD310, 3)) # <-----> Table 2
stopifnot(abs(relD2) < 0.006) # max is 0.00412
stopifnot(abs(relD310) < 1e-5 ) # max is 5.5953e-6

## Arguments as for Table 3 (p.152) of the reference :
aarg3 <- cbind(a = c( 10, 10, 10, 15, 10, 12, 30, 35),
                 b = c( 5, 10, 30, 20, 5, 17, 30, 30),
                 ncp=c( 20, 54, 80,120, 55, 64,140, 20),
                 x = c(644,700,780,760,795,560,800,670)/1000)
pnbA3 <- apply(aarg3, 1, function(aa) do.call(pnbtAppr2, as.list(aa)))
pnA310<- apply(aarg3, 1, function(aa) do.call(pnbtAS310, as.list(aa)))
aar3 <- aarg3; dimnames(aar3)[[2]] <- c(paste0("shape", 1:2), "ncp", "q")
pnbR <- apply(aar3, 1, function(aa) do.call(pbta, as.list(aa)))
range(relD2 <- 1 - pnbA3 /pnbR)
range(relD310 <- 1 - pnA310/pnbR)
cbind(aarg3, pnbA3, pnA310, pnbR,
      relD2 = signif(relD2, 3), relD310 = signif(relD310, 3)) # <-----> Table 3
stopifnot(abs(relD2) < 0.09) # max is 0.06337
stopifnot(abs(relD310) < 1e-4) # max is 3.898e-5

```

pnchi1sq

(Probabilities of Non-Central Chi-squared Distribution for Special Cases)

Description

Computes probabilities for the non-central chi-squared distribution, in special cases, currently for $df = 1$ and $df = 3$, using ‘exact’ formulas only involving the standard normal (Gaussian) cdf $\Phi()$ and its derivative $\phi()$, i.e., R’s **pnorm()** and **dnorm()**.

Usage

```

pnchi1sq(q, ncp = 0, lower.tail = TRUE, log.p = FALSE, epsS = .01)
pnchi3sq(q, ncp = 0, lower.tail = TRUE, log.p = FALSE, epsS = .04)

```

Arguments

<code>q</code>	number ('quantile', i.e., abscissa value.)
<code>ncp</code>	non-centrality parameter δ ;
<code>lower.tail, log.p</code>	logical, see, e.g., pchisq() .
<code>epsS</code>	small number, determining where to switch from the "small case" to the regular case, namely by defining <code>small <- sqrt(q/ncp) <= epsS</code> .

Details

In the "small case" (`epsS` above), the direct formulas suffer from cancellation, and we use Taylor series expansions in $s := \sqrt{q}$, which in turn use "probabilists" Hermite polynomials $He_n(x)$.

The default values `epsS` have currently been determined by experiments as those in the 'Examples' below.

Value

a numeric vector "like" `q+ncp`, i.e., recycled to common length.

Author(s)

Martin Maechler, notably the Taylor approximations in the "small" cases.

References

Johnson et al.(1995), see 'References' in [pchisqPearson](#).

https://en.wikipedia.org/wiki/Hermite_polynomials

See Also

[pchisq](#), the (simple and R-like) approximations, such as [pnchisqPearson](#) and the wienergerm approximations, [pchisqW\(\)](#) etc.

Examples

```
qq <- seq(9500, 10500, length=1000)
m1 <- cbind(pch = pchisq (qq, df=1, ncp = 10000),
             p1 = pnchi1sq(qq, ncp = 10000))
matplot(qq, m1, type = "l"); abline(h=0:1, v=10000+1, lty=3)
all.equal(m1[, "p1"], m1[, "pch"], tol=0) # for now, 2.37e-12

m3 <- cbind(pch = pchisq (qq, df=3, ncp = 10000),
             p3 = pnchi3sq(qq, ncp = 10000))
matplot(qq, m3, type = "l"); abline(h=0:1, v=10000+3, lty=3)
all.equal(m3[, "p3"], m3[, "pch"], tol=0) # for now, 1.88e-12

stopifnot(exprs = {
  all.equal(m1[, "p1"], m1[, "pch"], tol=1e-10)
  all.equal(m3[, "p3"], m3[, "pch"], tol=1e-10)
})

### Very small 'x' i.e., 'q' would lead to cancellation: -----
```

```

##  df = 1 ----

qS <- c(0, 2^seq(-40,4, by=1/16))
m1s <- cbind(pch = pnchi1sq(qS, df=1, ncp = 1)
              , p1.0= pnchi1sq(qS,           ncp = 1, epsS = 0)
              , p1.4= pnchi1sq(qS,           ncp = 1, epsS = 1e-4)
              , p1.3= pnchi1sq(qS,           ncp = 1, epsS = 1e-3)
              , p1.2= pnchi1sq(qS,           ncp = 1, epsS = 1e-2)
            )
cols <- adjustcolor(1:5, 1/2); lws <- seq(4,2, by = -1/2)
abl.leg <- function(x.leg = "topright", epsS = 10^{-(4:2)}, legend = NULL)
{
  abline(h = .Machine$double.eps, v = epsS^2,
         lty = c(2,3,3,3), col= adjustcolor(1, 1/2))
  if(is.null(legend))
    legend <- c(quote(epsS == 0), as.expression(lapply(epsS,
              function(K) substitute(epsS == KK,
                list(KK = formatC(K, w=1))))))
  legend(x.leg, legend, lty=1:4, col=cols, lwd=lws, bty="n")
}
matplot(qS, m1s, type = "l", log="y" , col=cols, lwd=lws)
matplot(qS, m1s, type = "l", log="xy", col=cols, lwd=lws) ; abl.leg("right")
## ===== "Errors" =====
## Absolute: -----
matplot(qS,      m1s[,1] - m1s[,-1] , type = "l", log="x" , col=cols, lwd=lws)
matplot(qS, abs(m1s[,1] - m1s[,-1]), type = "l", log="xy", col=cols, lwd=lws)
abl.leg("bottomright")
## Relative: -----
matplot(qS,      1 - m1s[,-1]/m1s[,1] , type = "l", log="x", col=cols, lwd=lws)
abl.leg()
matplot(qS, abs(1 - m1s[,-1]/m1s[,1]), type = "l", log="xy", col=cols, lwd=lws)
abl.leg()

##  df = 3 ----- %% FIXME: the 'small' case is clearly wrong <<<

qS <- c(0, 2^seq(-40,4, by=1/16))
ee <- c(1e-3, 1e-2, .04)
m3s <- cbind(pch = pnchi3sq(qS, df=3, ncp = 1)
              , p1.0= pnchi3sq(qS,           ncp = 1, epsS = 0)
              , p1.3= pnchi3sq(qS,           ncp = 1, epsS = ee[1])
              , p1.2= pnchi3sq(qS,           ncp = 1, epsS = ee[2])
              , p1.1= pnchi3sq(qS,           ncp = 1, epsS = ee[3])
            )
matplot(qS, m3s, type = "l", log="y" , col=cols, lwd=lws)
matplot(qS, m3s, type = "l", log="xy", col=cols, lwd=lws); abl.leg("right", ee)
## ===== "Errors" =====
## Absolute: -----
matplot(qS,      m3s[,1] - m3s[,-1] , type = "l", log="x" , col=cols, lwd=lws)
matplot(qS, abs(m3s[,1] - m3s[,-1]), type = "l", log="xy", col=cols, lwd=lws)
abl.leg("right", ee)
## Relative: -----
matplot(qS,      1 - m3s[,-1]/m3s[,1] , type = "l", log="x", col=cols, lwd=lws)
abl.leg(, ee)
matplot(qS, abs(1 - m3s[,-1]/m3s[,1]), type = "l", log="xy", col=cols, lwd=lws)
abl.leg(, ee)

```

Description

Compute (approximate) probabilities for the non-central chi-squared distribution.

The non-central chi-squared distribution with $df = n$ degrees of freedom and non-centrality parameter $ncp = \lambda$ has density

$$f(x) = f_{n,\lambda}(x) = e^{-\lambda/2} \sum_{r=0}^{\infty} \frac{(\lambda/2)^r}{r!} f_{n+2r}(x)$$

for $x \geq 0$; for more, see R's help page for [pchisq](#).

- R's own historical and current versions, but with more tuning parameters;

Historical relatively simple approximations listed in Johnson, Kotz, and Balakrishnan (1995):

- Patnaik(1949)'s approximation to the non-central via central chi-squared. Is also the formula 26.4.27 in Abramowitz & Stegun, p.942. Johnson et al mention that the approximation error is $O(1/\sqrt{(\lambda)})$ for $\lambda \rightarrow \infty$.
- Pearson(1959) is using 3 moments instead of 2 as Patnaik (to approximate via a central chi-squared), and therefore better than Patnaik for the right tail; further (in Johnson et al.), the approximation error is $O(1/\lambda)$ for $\lambda \rightarrow \infty$.
- Abdel-Aty(1954)'s “first approximation” based on Wilson-Hilferty via Gaussian ([pnorm](#)) probabilities, is partly *wrongly* cited in Johnson et al., p.463, eq.(29.61a).
- Bol'shev and Kuznetsov (1963) concentrate on the case of **small** $ncp \lambda$ and provide an “approximation” via *central* chi-squared with the same degrees of freedom df , but a modified q ('x'); the approximation has error $O(\lambda^3)$ for $\lambda \rightarrow 0$ and is from Johnson et al., p.465, eq.(29.62) and (29.63).
- Sankaran(1959, 1963) proposes several further approximations base on Gaussian probabilities, according to Johnson et al., p.463. [pnchisqSankaran_d\(\)](#) implements its formula (29.61d).

[pnchisq\(\)](#): an R implementation of R's own C [pnchisq_raw\(\)](#), but almost only up to Feb.27, 2004, long before the `log.p=TRUE` addition there, including *logspace arithmetic* in April 2014, its finish on 2015-09-01. Currently for historical reference only.

[pnchisqV\(\)](#): a [Vectorize\(\)](#)d [pnchisq](#).

[pnchisqRC\(\)](#): R's C implementation as of Aug.2019; but with many more options. Currently extreme cases tend to hang on Winbuilder (?)

[pnchisqIT](#):

[pnchisqTerms](#):

[pnchisqT93](#): pure R implementations of approximations when both q and ncp are large, by Temme(1993), from Johnson et al., p.467, formulas (29.71a), and (29.71b), using auxiliary functions [pnchisqT93a\(\)](#) and [pnchisqT93b\(\)](#) respectively, with adapted formulas for the `log.p=TRUE` cases.

[pnchisq_ss\(\)](#):

`ss`:

`ss2`:

`ss2.:`

Usage

```

pnchisq      (q, df, ncp = 0, lower.tail = TRUE,
               cutOffncp = 80, itSimple = 110, errmax = 1e-12, reltol = 1e-11,
               maxit = 10* 10000, verbose = 0, xLrg.sigma = 5)
pnchisqV(x, ..., verbose = 0)

pnchisqRC    (q, df, ncp = 0, lower.tail = TRUE, log.p = FALSE,
               no2nd.call = FALSE,
               cutOffncp = 80, small.ncp.logspace = small.ncp.logspaceR2015,
               itSimple = 110, errmax = 1e-12,
               reltol = 8 * .Machine$double.eps, epsS = reltol/2, maxit = 1e6,
               verbose = FALSE)
pnchisqAbdelAty (q, df, ncp = 0, lower.tail = TRUE, log.p = FALSE)
pnchisqBolkuz  (q, df, ncp = 0, lower.tail = TRUE, log.p = FALSE)
pnchisqPatnaik  (q, df, ncp = 0, lower.tail = TRUE, log.p = FALSE)
pnchisqPearson  (q, df, ncp = 0, lower.tail = TRUE, log.p = FALSE)
pnchisqSankaran_d(q, df, ncp = 0, lower.tail = TRUE, log.p = FALSE)
pnchisq_ss       (x, df, ncp = 0, lower.tail = TRUE, log.p = FALSE, i.max = 10000)
pnchisqTerms     (x, df, ncp,      lower.tail = TRUE, i.max = 1000)

pnchisqT93   (q, df, ncp, lower.tail = TRUE, log.p = FALSE, use.a = q > ncp)
pnchisqT93.a(q, df, ncp, lower.tail = TRUE, log.p = FALSE)
pnchisqT93.b(q, df, ncp, lower.tail = TRUE, log.p = FALSE)

ss   (x, df, ncp, i.max = 10000, useLv = !(expMin < -lambda && 1/lambda < expMax))
ss2  (x, df, ncp, i.max = 10000, eps = .Machine$double.eps)
ss2. (q, df, ncp = 0, errmax = 1e-12, reltol = 2 * .Machine$double.eps,
      maxit = 1e+05, eps = reltol, verbose = FALSE)

```

Arguments

<code>x</code>	numeric vector (of ‘quantiles’, i.e., abscissa values).
<code>q</code>	number (‘quantile’, i.e., abscissa value.)
<code>df</code>	degrees of freedom > 0 , maybe non-integer.
<code>ncp</code>	non-centrality parameter δ ;
<code>lower.tail, log.p</code>	logical, see, e.g., pchisq() .
<code>i.max</code>	number of terms in evaluation ...
<code>use.a</code>	logical vector for Temmle pnchisqT93*() formulas, indicating to use formula ‘a’ over ‘b’. The default is as recommended in the references, but they did not take into account <code>log.p = TRUE</code> situations.
<code>cutOffncp</code>	a positive number, the cutoff value for ncp...
<code>itSimple</code>	...
<code>errmax</code>	absolute error tolerance.
<code>reltol</code>	convergence tolerance for <i>relative</i> error.
<code>maxit</code>	maximal number of iterations.
<code>xLrg.sigma</code>	positive number ...

no2nd.call	logical indicating if a 2nd call is made to the internal function
small.ncp.logspace	logical vector or function , indicating if the logspace computations for “small” ncp (defined to fulfill ncp < cutOffncp !).
epsS	small positive number, the convergence tolerance of the ‘simple’ iterations...
verbose	logical or integer specifying if or how much the algorithm progress should be monitored.
...	further arguments passed from pnchisqV() to pnchisq().
useLv	logical indicating if logarithmic scale should be used for λ computations.
eps	convergence tolerance, a positive number.

Details

pnchisq_ss() uses $si \leftarrow ss(x, df, \dots)$ to get the series terms, and returns $2 * dchisq(x, df = df + 2) * sum(si$s)$.
 ss() computes the terms needed for the expansion used in pnchisq_ss().
 ss2() computes some simple “statistics” about ss(...).

Value

ss() returns a list with 3 components

s	the series
i1	location (in s[]) of the first change from 0 to positive.
max	(first) location of the maximal value in the series (i.e., which.max(s)).

Author(s)

Martin Maechler, from May 1999; starting from a post to the S-news mailing list by Ranjan Maitra (@ math.umbc.edu) who showed a version of our pchisqAppr.0() thanking Jim Stapleton for providing it.

References

- Johnson, N.L., Kotz, S. and Balakrishnan, N. (1995) Continuous Univariate Distributions Vol~2, 2nd ed.; Wiley.
 Chapter 29 *Noncentral χ^2 -Distributions*; notably Section 8 *Approximations*, p.461 ff.
 Abramowitz, M. and Stegun, I. A. (1972) *Handbook of Mathematical Functions*. New York: Dover.
https://en.wikipedia.org/wiki/Abramowitz_and_Stegun

See Also

- pchisq** and the wienergerm approximations for it: **pchisqW()** etc.
r_pois() and its plot function, for an aspect of the series approximations we use in pnchisq_ss().

Examples

```

## set of quantiles to use :
qq <- c(.001, .005, .01, .05, (1:9)/10, 2^seq(0, 10, by= 0.5))
## Take "all interesting" pchisq-approximation from our pkg :
pkg <- "package:DPQ"
pnchNms <- c(paste0("pchisq", c("V", "W", "W.", "W.R")),
             ls(pkg, pattern = "^pnchisq"))
pnchNms <- pnchNms[!grepl("Terms$", pnchNms)]
pnchF <- sapply(pnchNms, get, envir = as.environment(pkg))
str(pnchF)
ncps <- c(0, 1/8, 1/2)
pnchR <- as.list(setNames(ncps, paste("ncp",ncps, sep="")))
for(i.n in seq_along(ncps)) {
  ncp <- ncps[i.n]
  pnF <- if(ncp == 0) pnchF[!grepl("chisqT93", pnchNms)] else pnchF
  pnchR[[i.n]] <- sapply(pnF, function(F)
    Vectorize(F, names(formals(F))[[1]])(qq, df = 3, ncp=ncp))
}
str(pnchR, max=2)

## A case where the non-central P[] should be improved :
## First, the central P[] which is close to exact -- choosing df=2 allows
## truly exact values: chi^2 = Exp(1) !
opal <- palette()
palette(c("black", "red", "green3", "blue", "cyan", "magenta", "gold3", "gray44"))
cR <- curve(pchisq(x, df=2, lower.tail=FALSE, log.p=TRUE), 0, 4000, n=2001)
cRC <- curve(pnchisqRC(x, df=2, ncp=0, lower.tail=FALSE, log.p=TRUE),
              add=TRUE, col=adjustcolor(2,1/2), lwd=3, lty=2, n=2001)
cR0 <- curve(pchisq(x, df=2, ncp=0, lower.tail=FALSE, log.p=TRUE),
              add=TRUE, col=adjustcolor(3,1/2), lwd=4, n=2001)
## smart "named list" constructur :
list_ <- function(...)
  `names<-`((list(...), vapply(sys.call()[-1L], as.character, "")))
JKBfn <- list_(pnchisqPatnaik,
                 pnchisqPearson,
                 pnchisqAbdelAty,
                 pnchisqBolKuz,
                 pnchisqSankaran_d)
cl. <- setNames(adjustcolor(3+seq_along(JKBfn), 1/2), names(JKBfn))
lw. <- setNames(2+seq_along(JKBfn), names(JKBfn))
cR.JKB <- sapply(names(JKBfn), function(nmf) {
  curve(JKBfn[[nmf]](x, df=2, ncp=0, lower.tail=FALSE, log.p=TRUE),
        add=TRUE, col=cl.[[nmf]], lwd=lw.[[nmf]], lty=lw.[[nmf]], n=2001)
})
legend("bottomleft", c("pchisq", "pchisq.ncp=0", "pnchisqRC", names(JKBfn)),
       col=c(palette()[1], adjustcolor(2:3,1/2), cl.),
       lwd=c(1,3,4, lw.), lty=c(1,2,1, lw.))
palette(opal)# revert

all.equal(cRC, cR0, tol = 1e-15) # TRUE [for now]
## the problematic "jump" :
as.data.frame(cRC)[744:750,]
if(.Platform$OS.type == "unix")
  ## verbose=TRUE may reveal which branches of the algorithm are taken:
  pnchisqRC(1500, df=2, ncp=0, lower.tail=FALSE, log.p=TRUE, verbose=TRUE) #

```

```

## |--> -Inf currently

## The *two* principal cases (both lower.tail = {TRUE,FALSE} !), where
## "2nd call" happens *and* is currently beneficial :
dfs <- c(1:2, 5, 10, 20)
pL. <- pnchisqRC(.00001, df=dfs, ncp=0, log.p=TRUE, lower.tail=FALSE, verbose = TRUE)
pR. <- pnchisqRC( 100, df=dfs, ncp=0, log.p=TRUE,                                     verbose = TRUE)
## R's own non-central version (specifying 'ncp'):
pL0 <- pchisq (.00001, df=dfs, ncp=0, log.p=TRUE, lower.tail=FALSE)
pR0 <- pchisq ( 100, df=dfs, ncp=0, log.p=TRUE)
## R's *central* version, i.e., *not* specifying 'ncp' :
pL <- pchisq (.00001, df=dfs,           log.p=TRUE, lower.tail=FALSE)
pR <- pchisq ( 100, df=dfs,           log.p=TRUE)
cbind(pL., pL, relEc = signif(1-pL./pL, 3), relE0 = signif(1-pL./pL0, 3))
cbind(pR., pR, relEc = signif(1-pR./pR, 3), relE0 = signif(1-pR./pR0, 3))

```

Description

Functions implementing the two Wiener germ approximations to `pchisq()`, the (non-central) chi-squared distribution, and to `qchisq()` its inverse, the quantile function.

These have been proposed by Penev and Raykov (2000) who also listed a Fortran implementation.

In order to use them in numeric boundary cases, Martin Maechler has improved the original formulas.

Auxiliary functions:

`sW()`: The $s()$ as in the Wienergerm approximation, but using Taylor expansion when needed, i.e., $(x*ncp / df^2) \ll 1$.

`qs()`: ...

`z0()`: ...

`z.f()`: ...

`z.s()`: ...

.....

Usage

```

pchisqW. (q, df, ncp = 0, lower.tail = TRUE, log.p = FALSE,
          Fortran = TRUE, variant = c("s", "f"))
pchisqV (q, df, ncp = 0, lower.tail = TRUE, log.p = FALSE,
          Fortran = TRUE, variant = c("s", "f"))
pchisqW (q, df, ncp = 0, lower.tail = TRUE, log.p = FALSE, variant = c("s", "f"))
pchisqW.R(x, df, ncp = 0, lower.tail = TRUE, log.p = FALSE, variant = c("s", "f"),
           verbose = getOption("verbose"))

sW(x, df, ncp)

```

```
qs(x, df, ncp, f.s = sW(x, df, ncp), eps1 = 1/2, sMax = 1e+100)
z0(x, df, ncp)
z.f(x, df, ncp)
z.s(x, df, ncp, verbose = getOption("verbose"))
```

Arguments

<code>q,x</code>	vector of quantiles (main argument, see pchisq).
<code>df</code>	degrees of freedom (non-negative, but can be non-integer).
<code>ncp</code>	non-centrality parameter (non-negative).
<code>lower.tail,log.p</code>	<code>logical</code> , see pchisq .
<code>variant</code>	a <code>character</code> string, currently either "f" for the first or "s" for the second Wienergerm approximation in Penev & Raykov (2000).
<code>Fortran</code>	logical specifying if the Fortran or the C version should be used.
<code>verbose</code>	logical (or integer) indicating if or how much diagnostic output should be printed to the console during the computations.
<code>f.s</code>	a number must be a “version” of $s(x, df, ncp)$.
<code>eps1</code>	for <code>qs()</code> : use direct approximation instead of $h(1 - 1/s)$ for $s < \text{eps1}$.
<code>sMax</code>	for <code>qs()</code> : cutoff to switch the $h(\cdot)$ formula for $s > \text{sMax}$.

Details

....TODO... or write vignette

Value

all these functions return `numeric` vectors according to their arguments.

Note

The exact auxiliary function names etc, are still considered *provisional*; currently they are exported for easier documentation and use, but may well all disappear from the exported functions or even completely.

Author(s)

Martin Maechler, mostly end of Jan 2004

References

Penev, Spiridon and Raykov (2000) A Wiener Germ approximation of the noncentral chi square distribution and of its quantiles. *Computational Statistics* **15**, 219–228. doi: [10.1007/s001800000029](https://doi.org/10.1007/s001800000029)

Dinges, H. (1989) Special cases of second order Wiener germ approximations. *Probability Theory and Related Fields*, **83**, 5–57.

See Also

[pchisq](#), and other approximations for it: [pnchisq\(\)](#) etc.

Examples

```
#> see example(pnchisqAppr) which looks at all of the pchisq() approximating functions
```

pnt

Non-central t Probability Distribution - Algorithms and Approximations

Description

Compute different approximations for the non-central t-Distribution cumulative probability distribution function.

Usage

```
pntR      (t, df, ncp, lower.tail = TRUE, log.p = FALSE,
           use.pnorm = (df > 4e5 ||
                         ncp^2 > 2*log(2)*(-.Machine$double.min.exp)),
                         itrmax = 1000, errmax = 1e-12, verbose = TRUE)
pntR1     (t, df, ncp, lower.tail = TRUE, log.p = FALSE,
           use.pnorm = (df > 4e5 ||
                         ncp^2 > 2*log(2)*(-.Machine$double.min.exp)),
                         itrmax = 1000, errmax = 1e-12, verbose = TRUE)

pntP94    (t, df, ncp, lower.tail = TRUE, log.p = FALSE,
           itrmax = 1000, errmax = 1e-12, verbose = TRUE)
pntP94.1  (t, df, ncp, lower.tail = TRUE, log.p = FALSE,
           itrmax = 1000, errmax = 1e-12, verbose = TRUE)

pnt3150   (t, df, ncp, lower.tail = TRUE, log.p = FALSE, M = 1000, verbose = TRUE)
pnt3150.1 (t, df, ncp, lower.tail = TRUE, log.p = FALSE, M = 1000, verbose = TRUE)

pntLrg    (t, df, ncp, lower.tail = TRUE, log.p = FALSE)

pntJW39   (t, df, ncp, lower.tail = TRUE, log.p = FALSE)
pntJW39.0 (t, df, ncp, lower.tail = TRUE, log.p = FALSE)
```

Arguments

- t vector of quantiles (called q in `pt(...)`).
- df degrees of freedom (> 0 , maybe non-integer). `df = Inf` is allowed.
- ncp non-centrality parameter $\delta \geq 0$; If omitted, use the central t distribution.

log, log.p	logical; if TRUE, probabilities p are given as log(p).
lower.tail	logical; if TRUE (default), probabilities are $P[X \leq x]$, otherwise, $P[X > x]$.
use.pnorm	<code>logical</code> indicating if the <code>pnorm()</code> approximation of Abramowitz and Stegun (26.7.10) should be used, which is available as <code>pntLrg()</code> . The default corresponds to R <code>pt()</code> 's own behaviour (which is most probably suboptimal).
itrmax	number of iterations / terms.
errmax	convergence bound for the iterations.
verbose	<code>logical</code> or integer determining the amount of diagnostic print out to the console.
M	positive integer specifying the number of terms to use in the series.

Details

`pntR1()`: a pure R version of the (C level) code of R's own `pt()`, additionally giving more flexibility (via arguments `use.pnorm`, `itrmax`, `errmax` whose defaults here have been hard-coded in R's C code).

This implements an improved version of the AS 243 algorithm from Lenth(1989);

R's help on non-central `pt()` says: *This computes the lower tail only, so the upper tail suffers from cancellation and a warning will be given when this is likely to be significant.*

and (in 'Note:') *The code for non-zero ncp is principally intended to be used for moderate values of ncp: it will not be highly accurate, especially in the tails, for large values.*

`pntR()`: the `Vectorize()`d version of `pntR1()`.

`pntP94()`, `pntP94.1()`: New versions of `pntR1()`, `pntR()`; using the Posten (1994) algorithm.
`pntP94()` is the `Vectorize()`d version of `pntP94.1()`.

`pnt3150()`, `pnt3150.1()`: Simple inefficient but hopefully correct version of `pntP94..()` This is really a direct implementation of formula (31.50), p.532 of Johnson, Kotz and Balakrishnan (1995)

`pntLrg()`: provides the `pnorm()` approximation (to the non-central t) from Abramowitz and Stegun (26.7.10), p.949; which should be employed only for *large* df and/or ncp.

`pntJW39.0()`: use the Jennett & Welch (1939) approximation see Johnson et al. (1995), p. 520, after (31.26a). This is still *fast* for huge ncp but has *wrong* asymptotic tail for $|t| \rightarrow \infty$. Crucially needs $b = b_chi(df)$.

`pntJW39()`: is an improved version of `pntJW39.0()`, using $1 - b = b_chi(df, one.minus=TRUE)$ to avoid cancellation when computing $1 - b^2$.

Value

a number for `pntJKBf1()` and `.pntJKBch1()`.

a numeric vector of the same length as the maximum of the lengths of `x`, `df`, `ncp` for `pntJKBf()` and `.pntJKBch()`.

Author(s)

Martin Maechler

References

- Johnson, N.L., Kotz, S. and Balakrishnan, N. (1995) Continuous Univariate Distributions Vol~2, 2nd ed.; Wiley.
 Chapter 31, Section 5 *Distribution Function*, p.514 ff
- Lenth, R. V. (1989). *Algorithm AS 243 — Cumulative distribution function of the non-central t distribution*, *Applied Statistics* **38**, 185–189.
- Abramowitz, M. and Stegun, I. A. (1972) *Handbook of Mathematical Functions*. New York: Dover. Formula (26.7.10), p.949

See Also

[pt](#), for R's version of non-central t probabilities.

Examples

```
tt <- seq(0, 10, len = 21)
ncp <- seq(0, 6, len = 31)
dt3R <- outer(tt, ncp, pt, , df = 3)
dt3JKB <- outer(tt, ncp, pntR, df = 3)# currently verbose
stopifnot(all.equal(dt3R, dt3JKB, tolerance = 4e-15))# 64-bit Lnx: 2.78e-16
```

ppoisson

Direct Computation of 'ppois()' Poisson Distribution Probabilities

Description

Direct computation and errors of [ppois](#) Poisson distribution probabilities.

Usage

```
ppoisD(q, lambda, all.from.0 = TRUE, verbose = 0L)
ppoisErr (lambda, ppFUN = ppoisD, iP = 1e-15,
          xM = qpois(iP, lambda=lambda, lower.tail=FALSE),
          verbose = FALSE)
```

Arguments

<code>q</code>	numeric vector of non-negative integer values, “quantiles” at which to evaluate ppois (q,la) and ppFUN(q,la).
<code>lambda</code>	positive parameter of the Poisson distribution, <code>lambda</code> = $\lambda = E[X] = \text{Var}[X]$ where $X \sim \text{Pois}(\lambda)$.
<code>all.from.0</code>	<code>logical</code> indicating if <code>q</code> is positive integer, and the probabilities should be computed for all quantile values of $0:q$.
<code>ppFUN</code>	alternative ppois evaluation, by default the direct summation of dpois (k,lambda).
<code>iP</code>	small number, <code>iP</code> << 1, used to construct the abscissa values <code>x</code> at which to evaluate and compare ppois() and <code>ppFUN()</code> , see <code>xM</code> :
<code>xM</code>	(specified instead of <code>iP</code>) the maximal <code>x</code> -value to be used, i.e., the values used will be $x < -0:iM$. The default, <code>qpois(1-iP,lambda = lambda)</code> is the upper tail <code>iP</code> -quantile of <code>Poi(lambda)</code> .
<code>verbose</code>	<code>integer</code> (≥ 0) or <code>logical</code> indicating if extra information should be printed.

Value

`ppoisD()` contains the poisson probabilities along `q`, i.e., is a numeric vector of length `length(q)`.

`re <- ppoisErr()` returns the relative “error” of `ppois(x0, lambda)` where `ppFUN(x0, lambda)` is assumed to be the truth and `x0` the “worst case”, i.e., the value (among `x`) with the largest such difference.

Additionally, `attr(re, "x0")` contains that value `x0`.

Author(s)

Martin Maechler, March 2004; 2019 ff

See Also

[ppois](#)

Examples

```
(lams <- outer(c(1,2,5), 10^(0:3)))# 10^4 is already slow!
system.time(e1 <- sapply(lams, ppoisErr))
e1 / .Machine$double.eps

## Try another 'ppFUN' :-----
## this relies on the fact that it's *only* used on an 'x' of the form 0:M :
ppD0 <- function(x, lambda, all.from.0=TRUE)
  cumsum(dpois(if(all.from.0) 0:x else x, lambda=lambda))
## and test it:
p0 <- ppD0 ( 1000, lambda=10)
p1 <- ppois(0:1000, lambda=10)
stopifnot(all.equal(p0,p1, tol=8*.Machine$double.eps))

system.time(p0.slow <- ppoisD(0:1000, lambda=10, all.from.0=FALSE))# not very slow, here
p0.1 <- ppoisD(1000, lambda=10)
if(requireNamespace("Rmpfr")) {
  ppoisMpfr <- function(x, lambda) cumsum(Rmpfr::dpois(x, lambda=lambda))
  p0.best <- ppoisMpfr(0:1000, lambda = Rmpfr::mpfr(10, precBits = 256))
  AllEq. <- Rmpfr::all.equal
  AllEq <- function(target, current, ...)
    AllEq.(target, current, ...
      formatFUN = function(x, ...) Rmpfr::format(x, digits = 9))
  print(AllEq(p0.best, p0, tol = 0)) # 2.06e-18
  print(AllEq(p0.best, p0.slow, tol = 0)) # the "worst" (4.44e-17)
  print(AllEq(p0.best, p0.1, tol = 0)) # 1.08e-18
}

## Now (with 'all.from.0 = TRUE', it is fast too):
p15     <- ppoisErr(2^13)
p15.0. <- ppoisErr(2^13, ppFUN = ppD0)
c(p15, p15.0.) / .Machine$double.eps # on Lnx 64b, see (-10 2.5), then (-2 -2)

## lapply(), so you see "x0" values :
str(e0. <- lapply(lams, ppoisErr, ppFUN = ppD0))

## The first version [called 'err.lambd0()' for years] used simple cumsum(dpois(..))
## NOTE: It is *still* much faster, as it relies on special x == 0:M relation
## Author: Martin Maechler, Date: 1 Mar 2004, 17:40
```

```
##  
e0 <- sapply(lams, function(lamb) ppoisErr(lamb, ppFUN = ppD0))  
all.equal(e1, e0) # typically TRUE, though small "random" differences:  
cbind(e1, e0) * 2^53 # on Lnx 64b, seeing integer values in {-24, .., 33}
```

qbetaAppr*Compute (Approximate) Quantiles of the Beta Distribution***Description**

Compute quantiles (inverse distribution values) for the beta distribution, using diverse approximations.

Usage

```
qbetaAppr.1(a, p, q, y = qnormUappr(a))  
qbetaAppr.2(a, p, q, lower.tail=TRUE, log.p=FALSE, logbeta = lbeta(p,q))  
qbetaAppr.3(a, p, q, lower.tail=TRUE, log.p=FALSE, logbeta = lbeta(p,q))  
qbetaAppr.4(a, p, q, y = qnormUappr(a),  
            verbose = getOption("verbose"))  
  
qbetaAppr (a, p, q, y = qnormUappr(a), logbeta= lbeta(p,q),  
            verbose = getOption("verbose") && length(a) == 1)  
  
qbeta.R (alpha, p, q,  
         lower.tail = TRUE, log.p = FALSE,  
         logbeta = lbeta(p,q),  
         low.bnd = 3e-308, up.bnd = 1-2.22e-16,  
         method = c("AS109", "Newton-log"),  
         tol.outer = 1e-15,  
         f.acu = function(a,p,q) max(1e-300, 10^{(-13 - 2.5/pp^2 - .5/a^2)}),  
         fpu = .Machine$double.xmin,  
         qnormU.fun = function(u, lu) qnormUappr(p=u, lp=lu)  
               , R.pre.2014 = FALSE  
               , verbose = getOption("verbose")  
               , non.finite.report = verbose  
         )
```

Arguments

- a, alpha vector of probabilities (otherwise, e.g., in **qbeta()**, called p).
- p, q the two shape parameters of the beta distribution; otherwise, e.g., in **qbeta()**, called shape1 and shape2.
- y an approximation to $\Phi^{-1}(1 - \alpha)$ (aka $z_{1-\alpha}$) where $\Phi(x)$ is the standard normal cumulative probability function and $\Phi^{-1}(x)$ its inverse, i.e., R's **qnorm(x)**.
- lower.tail, log.p logical, see, e.g., **qchisq()**; must have length 1.
- logbeta must be **lbeta(p, q)**; mainly an option to pass a value already computed.

verbose	logical or integer indicating if and how much “monitoring” information should be produced by the algorithm.
low.bnd, up.bnd	lower and upper bounds for ...TODO...
method	a string specifying the approximation method to be used.
tol.outer	the “outer loop” convergence tolerance; the default 1e-15 has been hardwired in R’s qbeta() .
f.acu	a function with arguments (a, p, q) ...TODO...
fpu	a very small positive number.
qnormU.fun	a function with arguments (u, lu) to compute “the same” as qnormUapp() , the upper standard normal quantile.
R.pre.2014	a logical ... TODO ...
non.finite.report	a logical indicating if during the “outer loop” refining iterations, if y becomes non finite and the iterations have to stop, it should be reported (before the current best value is returned).

Value

...

Author(s)

The R Core Team for the C version in R’s sources; Martin Maechler for the R port.

See Also[qbета.](#)**Examples**

```

qbета.R(0.6, 2, 3) # 0.4445
qbета.R(0.6, 2, 3) - qbета(0.6, 2,3) # almost 0

qbетаRV <- Vectorize(qбeta.R, "alpha") # now can use
curve(qбetaRV(x, 1.5, 2.5))
curve(qбeta (x, 1.5, 2.5), add=TRUE, lwd = 3, col = adjustcolor("red", 1/2))

## an example of disagreement (and doubt, as borderline, close to underflow):
qbета.R(0.5078, .01, 5) # -> 2.77558e-15 # but
qbета (0.5078, .01, 5) # -> 1.776357e-15 now gives 4.651188e-31 !!!
qbета (0.5078, .01, 5, ncp=0) # also gives 4.651188e-31

```

qchisqAppr*Compute Approximate Quantiles of the Chi-Squared Distribution***Description**

Compute quantiles (inverse distribution values) for the chi-squared distribution. using Johnson,Kotz,...
.....TODO.....

Usage

```
qchisqKG (p, df, lower.tail = TRUE, log.p = FALSE)
qchisqWH (p, df, lower.tail = TRUE, log.p = FALSE)
qchisqAppr (p, df, lower.tail = TRUE, log.p = FALSE, tol = 5e-7)
qchisqAppr.R(p, df, lower.tail = TRUE, log.p = FALSE, tol = 5e-07,
  maxit = 1000, verbose = getOption("verbose"), kind = NULL)
```

Arguments

p	vector of probabilities.
df	degrees of freedom > 0, maybe non-integer; must have length 1.
lower.tail, log.p	logical, see, e.g., qchisq() ; must have length 1.
tol	non-negative number, the convergence tolerance
maxit	the maximal number of iterations
verbose	logical indicating if the algorithm should produce “monitoring” information.
kind	the <i>kind</i> of approximation; if NULL , the default, the approximation chosen depends on the arguments; notably it is chosen separately for each p . Otherwise, it must be a character string. The main approximations are Wilson-Hilferty versions, when the string contains "WH". More specifically, it must be one of the strings "chi.small" particularly useful for small chi-squared values p ;... ... "WH" "p1WH" "WHchk" "df.small" particularly useful for small degrees of freedom df

Value

...

Author(s)

Martin Maechler

See Also

[qchisq](#). Further, our approximations to the *non-central* chi-squared quantiles, [qnchisqAppr](#)

Examples

```
## TODO
```

qgammaAppr

Compute (Approximate) Quantiles of the Gamma Distribution

Description

Compute approximations to the quantile (i.e., inverse cumulative) function of the Gamma distribution.

Usage

```
qgammaAppr(p, shape, lower.tail = TRUE, log.p = FALSE,
           tol = 5e-07)
qgamma.R(p, alpha, scale = 1, lower.tail = TRUE, log.p = FALSE,
          EPS1 = 0.01, EPS2 = 5e-07, epsN = 1e-15, maxit = 1000,
          pMin = 1e-100, pMax = (1 - 1e-14),
          verbose = getOption("verbose"))

qgammaApprKG(p, shape, lower.tail = TRUE, log.p = FALSE)

qgammaApprSmallP(p, shape, lower.tail = TRUE, log.p = FALSE)
```

Arguments

p	numeric vector (possibly log tranformed) probabilities.
shape, alpha	shape parameter, non-negative.
scale	scale parameter, non-negative, see qgamma .
lower.tail, log.p	logical, see, e.g., qgamma() ; must have length 1.
tol	tolerance of maximal approximation error.
EPS1	small positive number. ...
EPS2	small positive number. ...
epsN	small positive number. ...
maxit	maximal number of iterations. ...
pMin, pMax	boundaries for p. ...
verbose	logical indicating if the algorithm should produce “monitoring” information.

Details

`qgammaApprSmallP(p, a)` should be a good approximation in the following situation when both `p` and `shape = alpha =: a` are small :

If we look at Abramowitz&Stegun $\gamma(a, x) = x^{-a} * P(a, x)$ and its series $g * (a, x) = 1/\gamma(a) * (1/a - 1/(a+1) * x + \dots)$,

then the first order approximation $P(a, x) = x^a * g * (a, x) = x^a / \gamma(a+1)$ and hence its inverse $x = qgamma(p, a) = (p * \gamma(a+1))^{(1/a)}$ should be good as soon as $1/a \gg 1/(a+1) * x$

$$\Leftrightarrow x \ll (a+1)/a = (1 + 1/a)$$

$$\Leftrightarrow x < \text{eps}^{(a+1)/a}$$

$$\Leftrightarrow \log(x) < \log(\text{eps}) + \log((a+1)/a) = \log(\text{eps}) + \log((a+1)/a) \sim -36 - \log(a) \text{ where } \log(x) \approx \log(p * \gamma(a+1)) / a = (\log(p) + \text{lgamma1p}(a)) / a$$

such that the above

$$\Leftrightarrow (\log(p) + \text{lgamma1p}(a)) / a < \log(\text{eps}) + \log((a+1)/a)$$

$$\Leftrightarrow \log(p) + \text{lgamma1p}(a) < a * (-\log(a) + \log(\text{eps}) + \log1p(a))$$

$$\Leftrightarrow \log(p) < a * (-\log(a) + \log(\text{eps}) + \log1p(a)) - \text{lgamma1p}(a) =: \text{bnd}(a)$$

Note that `qgammaApprSmallP()` indeed also builds on `lgamma1p()`.

. `qgammaApprBnd(a)` provides this bound `bnd(a)`; it is simply `a * (logEps + log1p(a) - log(a)) - lgamma1p(a)`, where `logEps` is $\log(\epsilon) = \log(\text{eps})$ where `eps <- .Machine$double.eps`, i.e. typically (always?) $\logEps = \log \epsilon = -52 * \log(2) = -36.04365$.

Value

numeric

Author(s)

Martin Maechler

References

Abramowitz, M. and Stegun, I. A. (1972) *Handbook of Mathematical Functions*. New York: Dover.
https://en.wikipedia.org/wiki/Abramowitz_and_Stegun provides links to the full text which is in public domain.

See Also

`qgamma` for R's Gamma distribution functions.

Examples

```
## TODO : Move some of the curve()s from ../tests/qgamma-ex.R !!
```

qnchisqAppr	<i>Compute Approximate Quantiles of Noncentral Chi-Squared Distribution</i>
-------------	---

Description

Compute quantiles (inverse distribution values) for the *non-central* chi-squared distribution.
..... using Johnson,Kotz, and other approximations

Usage

```

qchisqAppr.0 (p, df, ncp = 0, lower.tail = TRUE, log.p = FALSE)
qchisqAppr.1 (p, df, ncp = 0, lower.tail = TRUE, log.p = FALSE)
qchisqAppr.2 (p, df, ncp = 0, lower.tail = TRUE, log.p = FALSE)
qchisqAppr.3 (p, df, ncp = 0, lower.tail = TRUE, log.p = FALSE)
qchisqApprCF1(p, df, ncp = 0, lower.tail = TRUE, log.p = FALSE)
qchisqApprCF2(p, df, ncp = 0, lower.tail = TRUE, log.p = FALSE)

qchisqCappr.2 (p, df, ncp = 0, lower.tail = TRUE, log.p = FALSE)
qchisqN      (p, df, ncp = 0, qIni = qchisqAppr.0, ...)

qnchisqAbdelAty (p, df, ncp = 0, lower.tail = TRUE, log.p = FALSE)
qnchisqBolKuz   (p, df, ncp = 0, lower.tail = TRUE, log.p = FALSE)
qnchisqPatnaik  (p, df, ncp = 0, lower.tail = TRUE, log.p = FALSE)
qnchisqPearson  (p, df, ncp = 0, lower.tail = TRUE, log.p = FALSE)
qnchisqSankaran_d(p, df, ncp = 0, lower.tail = TRUE, log.p = FALSE)

```

Arguments

p	vector of probabilities.
df	degrees of freedom > 0, maybe non-integer.
ncp	non-centrality parameter δ ;
lower.tail, log.p	logical, see, e.g., qchisq() .
qIni	a function that computes an approximate noncentral chi-squared quantile as starting value x_0 for the Newton algorithm newton() .
...	further arguments to newton() , notably eps or maxiter.

Details

Compute (approximate) quantiles, using approximations analogous to those for the probabilities, see [pnchisqPearson](#).

qchisqAppr.0(): ...TODO...
qchisqAppr.1(): ...TODO...
qchisqAppr.2(): ...TODO...
qchisqAppr.3(): ...TODO...
qchisqApprCF1(): ...TODO...

qchisqApprCF2(): ...TODO...
qchisqCappr.2(): ...TODO...
qchisqN(): Uses Newton iterations with [pchisq\(\)](#) and [dchisq\(\)](#) to determine [qchisq\(. \)](#) values.
qnchisqAbdelAty(): ...TODO...
qnchisqBolKuz(): ...TODO...
qnchisqPatnaik(): ...TODO...
qnchisqPearson(): ...TODO...
qnchisqSankaran_d(): ...TODO...

Value

[numeric](#) vectors of (noncentral) chi-squared quantiles, corresponding to probabilities p.

Author(s)

Martin Maechler, from May 1999; starting from a post to the S-news mailing list by Ranjan Maitra (@ math.umbc.edu) who showed a version of our [qchisqAppr.0\(\)](#) thanking Jim Stapleton for providing it.

References

Johnson, N.L., Kotz, S. and Balakrishnan, N. (1995) Continuous Univariate Distributions Vol~2, 2nd ed.; Wiley.
 Chapter 29 *Noncentral χ^2 -Distributions*; notably Section 8 *Approximations*, p.461 ff.

See Also

[qchisq](#).

Examples

```
pp <- c(.001, .005, .01, .05, (1:9)/10, .95, .99, .995, .999)
pkg <- "package:DPQ"
qnchNms <- c(paste0("qchisqAppr.",0:3), paste0("qchisqApprCF",1:2),
              "qchisqN", "qchisqCappr.2", ls(pkg, pattern = "^qnchisq"))
qnchF <- sapply(qnchNms, get, envir = as.environment(pkg))
for(ncp in c(0, 1/8, 1/2)) {
  cat("\n~~~~~\nnncp: ", ncp, "\n=====\\n")
  print(sapply(qnchF, function(F) Vectorize(F, "p")(pp, df = 3, ncp=ncp)))
}

## Bug: qnchisqSankaran_d() has numeric overflow problems for large df:
qnchisqSankaran_d(pp, df=1e200, ncp = 100)

## One current (2019-08) R bug: Noncentral chi-squared quantiles on *LOG SCALE*

## a) left/lower tail : -----
qs <- 2^seq(0,11, by=1/16)
pql <- pchisq(qs, df=5, ncp=1, log.p=TRUE)
plot(qs, -pql, type="l", log="xy") # + expected warning on log(0) -- all fine
qpql <- qchisq(pql, df=5, ncp=1, log.p=TRUE) # severe overflow :
qm <- cbind(qs, pql, qchisq=qpql
, qchA.0 = qchisqAppr.0 (pql, df=5, ncp=1, log.p=TRUE)
```

```

, qchA.1 = qchisqAppr.1 (pqL, df=5, ncp=1, log.p=TRUE)
, qchA.2 = qchisqAppr.2 (pqL, df=5, ncp=1, log.p=TRUE)
, qchA.3 = qchisqAppr.3 (pqL, df=5, ncp=1, log.p=TRUE)
, qchACF1= qchisqApprCF1(pqL, df=5, ncp=1, log.p=TRUE)
, qchACF2= qchisqApprCF2(pqL, df=5, ncp=1, log.p=TRUE)
, qchCa.2= qchisqCappr.2(pqL, df=5, ncp=1, log.p=TRUE)
, qnPatnaik   = qnchisqPatnaik   (pqL, df=5, ncp=1, log.p=TRUE)
, qnAbdelAty = qnchisqAbdelAty (pqL, df=5, ncp=1, log.p=TRUE)
, qnBolKuz   = qnchisqBolKuz   (pqL, df=5, ncp=1, log.p=TRUE)
, qnPearson  = qnchisqPearson  (pqL, df=5, ncp=1, log.p=TRUE)
, qnSankaran_d= qnchisqSankaran_d(pqL, df=5, ncp=1, log.p=TRUE)
)

round(qm[ qs %in% 2^(0:11) , -2])
#=> Approximations don't overflow but are not good enough

## b) right/upper tail , larger ncp -----
qS <- 2^seq(-3, 3, by=1/8)
pqLu <- pchisq(qS, df=5, ncp=100, log.p=TRUE, lower.tail=FALSE)
## using "the alternative" (here is currently identical):
identical(pqLu, (pqLu.<- log1p(-pchisq(qS, df=5, ncp=100)))) # here TRUE
plot (qS, -pqLu, type="l", log="xy") # fine
qqLu <- qchisq(pqLu, df=5, ncp=100, log.p=TRUE, lower.tail=FALSE)
cbind(qS, pqLu, pqLu, qqLu)# severe underflow
qchMat <- cbind(qchisq = qqLu
, qchA.0 = qchisqAppr.0 (pqLu, df=5, ncp=100, log.p=TRUE, lower.tail=FALSE)
, qchA.1 = qchisqAppr.1 (pqLu, df=5, ncp=100, log.p=TRUE, lower.tail=FALSE)
, qchA.2 = qchisqAppr.2 (pqLu, df=5, ncp=100, log.p=TRUE, lower.tail=FALSE)
, qchA.3 = qchisqAppr.3 (pqLu, df=5, ncp=100, log.p=TRUE, lower.tail=FALSE)
, qchACF1= qchisqApprCF1(pqLu, df=5, ncp=100, log.p=TRUE, lower.tail=FALSE)
, qchACF2= qchisqApprCF2(pqLu, df=5, ncp=100, log.p=TRUE, lower.tail=FALSE)
, qchCa.2= qchisqCappr.2(pqLu, df=5, ncp=100, log.p=TRUE, lower.tail=FALSE)
, qnPatnaik   = qnchisqPatnaik   (pqLu, df=5, ncp=100, log.p=TRUE, lower.tail=FALSE)
, qnAbdelAty = qnchisqAbdelAty (pqLu, df=5, ncp=100, log.p=TRUE, lower.tail=FALSE)
, qnBolKuz   = qnchisqBolKuz   (pqLu, df=5, ncp=100, log.p=TRUE, lower.tail=FALSE)
, qnPearson  = qnchisqPearson  (pqLu, df=5, ncp=100, log.p=TRUE, lower.tail=FALSE)
, qnSankaran_d= qnchisqSankaran_d(pqLu, df=5, ncp=100, log.p=TRUE, lower.tail=FALSE)
)
cbind(L2err <- sort(sqrt(colSums((qchMat - qS)^2))))
##--> "Sankaran_d", "CF1" and "CF2" are good here

plot (qS, qqLu, type = "b", log="x", lwd=2)
lines(qS, qS, col="gray", lty=2, lwd=3)
top3 <- names(L2err)[1:3]
use <- c("qchisq", top3)
matlines(qS, qchMat[, use]) # 3 of the approximations are "somewhat ok"
legend("topleft", c(use,"True"), bty="n", col=c(palette()[1:4], "gray"),
lty = c(1:4,2), lwd = c(2, 1,1,1, 3))

```

Description

Relatively simple approximations to the standard normal (aka “Gaussian”) quantiles, i.e., the inverse of the normal cumulative probability function.

qnormUappr() is a simple approximation to (the **upper tail**) standard normal quantiles, [qnorm\(\)](#).

Usage

```
qnormAppr(p)
qnormUappr(p, lp = .DT_Clog(p, lower.tail=lower.tail, log.p=log.p),
            lower.tail = FALSE, log.p = FALSE)
```

Arguments

<i>p</i>	numeric vector of probabilities, possibly transformed, depending on <i>log.p</i> . Does not need to be specified, if <i>lp</i> is instead.
<i>lp</i>	$\log(1 - p^*)$, assuming p^* is the <i>lower.tail=TRUE, log.p=FALSE</i> version of <i>p</i> . If passed as argument, it can be much more accurate than when computed from <i>p</i> by default.
<i>lower.tail</i>	logical; if TRUE (<i>not</i> the default here!), probabilities are $P[X \leq x]$, otherwise (by default) upper tail probabilities, $P[X > x]$.
<i>log.p</i>	logical; if TRUE, probabilities <i>p</i> are given as $\log(p)$ in argument <i>p</i> .

Details

qnormAppr(p) uses the simple 4 coefficient rational approximation to [qnorm\(p\)](#), to be used *only* for $p > 1/2$ in [qbta\(\)](#) computations, e.g., [qbta.R](#).

The relative error of this approximation is quite *asymmetric*: It is mainly < 0 .

qnormUappr(p) uses the same rational approximation directly for the **Upper tail** where it is relatively good, and for the lower tail via “swapping the tails”, so it is good there as well.

Value

numeric vector of (approximate) normal quantiles corresponding to probabilities *p*

Author(s)

Martin Maechler

See Also

[qnorm](#).

Examples

```
pp <- c(.001, .005, .01, .05, (1:9)/10, .95, .99, .995, .999)
z_p <- qnorm(pp)
(R <- cbind(pp, z_p, qA = qnormAppr(pp), qUA = qnormUappr(pp, lower.tail=TRUE)))
## Errors, absolute and relative:
cbind(pp, (relE <- cbind(
  errA = z_p - R[, "qA"],
  errUA = z_p - R[, "qUA"]),
  rE.A = 1 - R[, "qA"]/z_p,
  rE.UA = 1 - R[, "qUA"]/z_p)))

lp <- -c(1000, 500, 200, 100, 50, 20:10, seq(9.75, 0, by = -1/8))
qnormUappr(lp=lp) # 'p' need not be specified if 'lp' is
```

```
curve(qnorm(x, lower.tail=FALSE), n=1001)
curve(qnormUappr(x), add=TRUE,      n=1001, col = adjustcolor("red", 1/2))

curve(qnorm(x, lower.tail=FALSE) - qnormUappr(x), n=1001)
```

qtAppr*Compute Approximate Quantiles of Non-Central t Distribution*

Description

Compute quantiles (inverse distribution values) for the non-central t distribution. using Johnson,Kotz,.. p.521, formula (31.26 a) (31.26 b) & (31.26 c)

Note that [qt\(...,ncp=*\)](#) did not exist yet in 1999, when MM implemented [qtAppr\(\)](#).

Usage

```
qtAppr(p, df, ncp, lower.tail = TRUE, log.p = FALSE, method = c("a", "b", "c"))
```

Arguments

- | | |
|-------------------|--|
| p | vector of probabilities. |
| df | degrees of freedom > 0 , maybe non-integer. |
| ncp | non-centrality parameter δ ; |
| lower.tail, log.p | logical, see, e.g., qt() . |
| method | a string specifying the approximation method to be used. |

Value

...

Author(s)

Martin Maechler, 6 Feb 1999

See Also

[qt](#).

Examples

```
## TODO
```

r_pois*Compute Relative Size of i-th term of Poisson Distribution Series***Description**

Compute

$$r_\lambda(i) := (\lambda^i / i!)/e_{i-1}(\lambda),$$

where $\lambda = \text{lambda}$, and

$$e_n(x) := 1 + x + x^2/2! + \dots + x^n/n!$$

is the n -th partial sum of $\exp(x) = e^x$.Questions: As function of i

- Can this be put in a simple formula, or at least be well approximated for large λ and/or large i ?
- For which i ($:= i_m(\lambda)$) is it maximal?
- When does $r_\lambda(i)$ become smaller than $(f+2i-x)/x = a + b*i$?

NB: This is relevant in computations for non-central chi-squared (and similar non-central distribution functions) defined as weighted sum with “Poisson weights”.

Usage

```
r_pois(i, lambda)
r_pois_expr # the R expression() for the asymptotic branch of r_pois()

p1Rpois(lambda, iset = 1:(2*lambda), do.main = TRUE,
         log = 'xy', type = "o", cex = 0.4, col = c("red","blue"),
         do.eaxis = TRUE, sub10 = "10")
```

Arguments

i	integer ..
lambda	non-negative number ...
iset
do.main	logical specifying if a main title should be drawn via (main = r_pois_expr).
type	type of (line) plot, see lines .
log	string specifying if (and where) logarithmic scales should be used, see plot.default() .
cex	character expansion factor.
col	colors for the two curves.
do.eaxis	logical specifying if eaxis () (package sfsmisc) should be used.
sub10	argument for eaxis () (with a different default than the original).

Details

r_pois() is related to our series expansions and approximations for the non-central chi-squared; in particular

p1Rpois() simply produces a “nice” plot of **r_pois(ii,*)** vs **ii**.

Value

`r_pois()` returns a numeric vector $r_\lambda(i)$ values.
`r_pois_expr()` an [expression](#).

Author(s)

Martin Maechler, 20 Jan 2004

See Also

[dpois\(\)](#).

Examples

```
plRpois(12)
plRpois(120)
```

Index

- *Topic **arith**
 - `logspace.add`, 23
- *Topic **character**
 - `format01prec`, 16
- *Topic **distribution**
 - `b_chi`, 6
 - `dchisqApprox`, 9
 - `dgamma-utils`, 10
 - `dgamma.R`, 11
 - `dnt`, 12
 - `DPQ-package`, 2
 - `dtWV`, 14
 - `lgamma1p`, 19
 - `lssum`, 24
 - `pbetaRv1`, 31
 - `pnbeta`, 33
 - `pnchi1sq`, 35
 - `pnchisqAppr`, 38
 - `pnchisqWienergerm`, 42
 - `pnt`, 44
 - `ppoisson`, 46
 - `qbetaAppr`, 48
 - `qchisqAppr`, 50
 - `qgammaAppr`, 51
 - `qnchisqAppr`, 53
 - `qnormAppr`, 55
 - `qtAppr`, 57
 - `r_pois`, 58
- *Topic **hplot**
 - `pl2curves`, 32
- *Topic **math**
 - `algdiv`, 5
 - `b_chi`, 6
 - `dchisqApprox`, 9
 - `dgamma.R`, 11
 - `dnt`, 12
 - `DPQ-package`, 2
 - `dtWV`, 14
 - `lbeta`, 17
 - `lgamma1p`, 19
 - `log1mexp`, 20
 - `log1pmx`, 21
 - `logcf`, 22
- `newton`, 26
- `numer-utils`, 29
- `pbetaRv1`, 31
- `pnbeta`, 33
- `pnchi1sq`, 35
- `pnchisqWienergerm`, 42
- `pnt`, 44
- *Topic **package**
 - `DPQ-package`, 2
- *Topic **print**
 - `format01prec`, 16
 - `.Machine`, 30
 - `.dntJKBch (dnt)`, 12
 - `.dntJKBch1 (dnt)`, 12
 - `.qgammaApprBnd (qgammaAppr)`, 51
- `algdiv`, 5, 18
- `all`, 30
- `all.equal`, 13
- `all_mpfr (numer-utils)`, 29
- `any`, 30
- `any_mpfr (numer-utils)`, 29
- `b_chi`, 6, 45
- `b_chiAsymp (b_chi)`, 6
- `bd0 (dgamma-utils)`, 10
- `besselI`, 9
- `beta`, 6, 17, 18
- `betaI (lbeta)`, 17
- `c_dt (b_chi)`, 6
- `c_dtAsymp (b_chi)`, 6
- `c_pt (b_chi)`, 6
- `character`, 16, 43, 50
- `class`, 12, 29
- `curve`, 32, 33
- `dchisq`, 9, 10, 40, 54
- `dchisqApprox`, 9
- `dchisqAsym (dchisqApprox)`, 9
- `dgamma`, 10, 11
- `dgamma-utils`, 10
- `dgamma.R`, 11
- `dnchisqBessel (dchisqApprox)`, 9

dnchisqR (dchisqApprox), 9
 dnoncentchisq (dchisqApprox), 9
 dnorm, 35
 dnt, 12
 dntJKBF, 15
 dntJKBF (dnt), 12
 dntJKBF1 (dnt), 12
 dpois, 9, 11, 46, 59
 dpois_raw, 12
 dpois_raw (dgamma-utils), 10
 DPQ (DPQ-package), 2
 DPQ-package, 2
 dt, 6, 13, 15
 dtWV, 13, 14
 eaxis, 58
 expression, 59
 format, 16
 format.pval, 16
 format01prec, 16
 formatC, 16
 function, 16, 27, 32, 40, 49, 53
 g2 (pnchisqWienergerm), 42
 G_half (numer-utils), 29
 gamma, 6, 19
 gnt (pnchisqWienergerm), 42
 h (pnchisqWienergerm), 42
 h0 (pnchisqWienergerm), 42
 h1 (pnchisqWienergerm), 42
 h2 (pnchisqWienergerm), 42
 hnt (pnchisqWienergerm), 42
 integer, 46
 lb_chi0 (b_chi), 6
 lb_chi00 (b_chi), 6
 lb_chiAsymp (b_chi), 6
 lbeta, 17, 18, 48
 lbeta_asy (lbeta), 17
 lbetaI (lbeta), 17
 lbetaM (lbeta), 17
 lbetaMM (lbeta), 17
 legend, 32
 length, 23
 lgamma, 6, 19
 lgamma1p, 19, 22, 23, 52
 lgamma1p_series (lgamma1p), 19
 lines, 58
 list, 27
 log, 6, 19, 29
 log1mexp, 20, 20, 23
 log1p, 19, 21, 22
 log1pmx, 19, 21, 22, 23
 logcf, 21, 22, 22
 logical, 17, 24, 27, 30, 34, 39, 40, 43, 45, 46, 49, 58
 logQab_asy, 6
 logQab_asy (lbeta), 17
 logr (numer-utils), 29
 logspace.add, 23
 logspace.sub (logspace.add), 23
 lssum, 23, 24, 26
 lsum, 23, 24, 25
 M_cutoff (numer-utils), 29
 M_LN2 (numer-utils), 29
 M_minExp (numer-utils), 29
 max, 24
 mpfr, 12, 13
 newton, 26, 53
 numer-utils, 29
 numeric, 11, 23, 43, 54
 okLongDouble (numer-utils), 29
 pbeta, 5, 18, 19, 23, 31, 32, 34
 pbetaRv1, 31
 pchisq, 36, 38–40, 42, 43, 54
 pchisqV (pnchisqWienergerm), 42
 pchisqW, 36, 40
 pchisqW (pnchisqWienergerm), 42
 pgamma, 23
 pl2curves, 32
 plot.default, 58
 plRpois (r_pois), 58
 pnbeta, 33
 pnbetaAppr2 (pnbeta), 33
 pnbetaAppr2v1 (pnbeta), 33
 pnbetaAS310 (pnbeta), 33
 pnchi1sq, 35
 pnchi3sq (pnchi1sq), 35
 pnchisq, 38, 43
 pnchisq (pnchisqAppr), 38
 pnchisq_ss (pnchisqAppr), 38
 pnchisqAbdelAty (pnchisqAppr), 38
 pnchisqAppr, 38
 pnchisqBolKuz (pnchisqAppr), 38
 pnchisqIT (pnchisqAppr), 38
 pnchisqPatnaik (pnchisqAppr), 38
 pnchisqPearson, 36, 53
 pnchisqPearson (pnchisqAppr), 38
 pnchisqRC (pnchisqAppr), 38
 pnchisqSankaran_d (pnchisqAppr), 38

pnchisqT93 (pnchisqAppr), 38
 pnchisqTerms (pnchisqAppr), 38
 pnchisqV (pnchisqAppr), 38
 pnchisqWienergerm, 42
 pnorm, 35, 38, 45
 pnt, 44
 pnt3150 (pnt), 44
 pntChShP94 (pnt), 44
 pntJW39, 6, 8
 pntJW39 (pnt), 44
 pntLrg (pnt), 44
 pntP94 (pnt), 44
 pntR (pnt), 44
 pntR1 (pnt), 44
 ppois, 46, 47
 ppoisD (ppoisson), 46
 ppoisErr (ppoisson), 46
 ppoisson, 46
 pt, 8, 44–46
 Qab_terms (lbeta), 17
 qbeta, 18, 48, 49, 56
 qbeta.R, 56
 qbeta.R (qbetaAppr), 48
 qbetaAppr, 48
 qchisq, 42, 48, 50, 53, 54
 qchisqAppr, 50
 qchisqAppr.0 (qnchisqAppr), 53
 qchisqAppr.1 (qnchisqAppr), 53
 qchisqAppr.2 (qnchisqAppr), 53
 qchisqAppr.3 (qnchisqAppr), 53
 qchisqApprCF1 (qnchisqAppr), 53
 qchisqApprCF2 (qnchisqAppr), 53
 qchisqCappr.2 (qnchisqAppr), 53
 qchisqKG (qchisqAppr), 50
 qchisqN, 28
 qchisqN (qnchisqAppr), 53
 qchisqWH (qchisqAppr), 50
 qgamma, 51, 52
 qgamma.R (qgammaAppr), 51
 qgammaAppr, 51
 qgammaApprKG (qgammaAppr), 51
 qgammaApprSmallP (qgammaAppr), 51
 qnchisqAbdelAty (qnchisqAppr), 53
 qnchisqAppr, 50, 53
 qnchisqBolKuz (qnchisqAppr), 53
 qnchisqPatnaik (qnchisqAppr), 53
 qnchisqPearson (qnchisqAppr), 53
 qnchisqSankaran_d (qnchisqAppr), 53
 qnorm, 48, 56
 qnormAppr, 55
 qnormUappr, 49
 qnormUappr (qnormAppr), 55
 qs (pnchisqWienergerm), 42
 qt, 57
 qtAppr, 57
 r_pois, 40, 58
 r_pois_expr (r_pois), 58
 Rmpfr, 4
 scalefactor (pnchisqWienergerm), 42
 ss (pnchisqAppr), 38
 ss2 (pnchisqAppr), 38
 stirlerr (dgamma-utils), 10
 sum, 24
 sW (pnchisqWienergerm), 42
 title, 58
 TRUE, 29, 30
 uniroot, 28
 Vectorize, 12, 38, 45
 warning, 27
 which.max, 40
 z.f (pnchisqWienergerm), 42
 z.s (pnchisqWienergerm), 42
 z0 (pnchisqWienergerm), 42