Package ‘mhn’

May 27, 2026

Type Package

Title The Modified Half-Normal Distribution

Version 0.1.0

Description Provides density, distribution, quantile, and random generation functions for the Modified Half-Normal (MHN) distribution, along with moments, mode, and the Fox-Wright Psi function used as the normalizing constant. The MHN distribution arises as a conditional posterior in Bayesian MCMC and generalizes the half-normal, truncated normal, and square-root gamma distributions. Implements efficient sampling via the Sun, Kong & Pal (2023) <doi:10.1080/03610926.2021.1934700> algorithms and the Gao & Wang (2025) <doi:10.1080/03610918.2025.2524551> RTDR method.

License MIT + file LICENSE

Encoding UTF-8

RoxygenNote 7.3.3

Depends R (>= 4.0.0)

Imports stats, Rcpp (>= 1.0.7)

LinkingTo Rcpp, BH (>= 1.78.0-0)

SystemRequirements C++17

Suggests testthat (>= 3.0.0), knitr, rmarkdown, bench, Rmpfr

Config/testthat/edition 3

VignetteBuilder knitr

URL `https://github.com/t-momozaki/mhn`, `https://t-momozaki.github.io/mhn/`

BugReports `https://github.com/t-momozaki/mhn/issues`

NeedsCompilation yes

Author Tomotaka Momozaki [aut, cre]

Maintainer Tomotaka Momozaki <momozaki.stat@gmail.com>

Repository CRAN

Date/Publication 2026-05-27 08:30:02 UTC

Contents

---

dmhn *Density of the Modified Half-Normal Distribution*

Description

Computes the probability density function (or log-density) of the Modified Half-Normal (MHN) distribution with parameters `alpha`, `beta`, and `gamma`.

Usage

```
dmhn(x, alpha = 1, beta = 1, gamma = 0, log = FALSE)
```

Arguments

`x`	Numeric vector of evaluation points.
`alpha`	Shape parameter ($\alpha > 0$). Scalar or numeric vector. Default: 1.
`beta`	Scale parameter ($\beta > 0$). Scalar or numeric vector. Default: 1.
`gamma`	Location parameter ($\gamma \in R$). Scalar or numeric vector. Default: 0.
`log`	Logical; if `TRUE`, log-density is returned. Default: `FALSE`.

Details

The MHN density is

$$f(x \mid \alpha, \beta, \gamma) = \frac{2\beta^{\alpha/2} x^{\alpha-1} \exp(-\beta x^2 + \gamma x)}{\Psi[\alpha/2, \gamma/\sqrt{\beta}]} \quad (x > 0)$$

where $\Psi[a, z]$ is the Fox-Wright Psi function (Sun et al., 2023, Lemma 1a).

The default parameters `alpha = 1, beta = 1, gamma = 0` correspond to the half-normal distribution $\mathrm{HN}(1/\sqrt{2})$.

Special cases are detected and dispatched to closed-form solutions:

- $\gamma = 0$: sqrt-Gamma distribution

- $\alpha = 1$: truncated normal distribution

Computation is performed in log-space to avoid numerical underflow/overflow.

When any of `alpha`, `beta`, `gamma` is a vector, the density is evaluated element-wise. The Fox-Wright Ψ normalizing constant is recomputed only when consecutive elements present a different (α, β, γ) triple, so passing grouped parameters is significantly faster than calling `dmhn` inside an R loop.

Value

A numeric vector. The output length equals `max(length(x), length(alpha), length(beta), length(gamma))`; each input is recycled to that length following standard R recycling rules. For `x < 0`, the density is 0 (`-Inf` if `log = TRUE`).

References

Sun, J., Kong, M., & Pal, S. (2023). The Modified-Half-Normal distribution: Properties and an efficient sampling scheme. *Communications in Statistics - Theory and Methods*, 52(5), 1507–1536.

See Also

`mhn_mean`, `mhn_var`, `mhn_mode`

Examples

```
x <- seq(0, 5, length.out = 100)
plot(x, dmhn(x, alpha = 2, beta = 1, gamma = 1), type = "l")

# Log-density
dmhn(1, alpha = 2, beta = 1, gamma = 1, log = TRUE)
```

`mhn_kurtosis` *Excess Kurtosis of the Modified Half-Normal Distribution*

Description

Computes the excess kurtosis $\gamma_2 = E[(X-\mu)^4]/\sigma^4 - 3$ for $X \sim \mathrm{MHN}(\alpha, \beta, \gamma)$.

Usage

```
mhn_kurtosis(alpha, beta, gamma)
```

Arguments

`alpha`	Shape parameter ($\alpha > 0$).
`beta`	Scale parameter ($\beta > 0$).
`gamma`	Location parameter ($\gamma \in R$).

Details

Uses the moment recurrence (Sun et al., 2023, Lemma 2b) to compute raw moments up to fourth order, then converts to central moments.

Value

A numeric scalar.

References

Sun, J., Kong, M., & Pal, S. (2023). The Modified-Half-Normal distribution: Properties and an efficient sampling scheme. *Communications in Statistics - Theory and Methods*, 52(5), 1507–1536. (Lemma 2b)

See Also

mhn_skewness, mhn_mean

Examples

```
mhn_kurtosis(alpha = 2, beta = 1, gamma = 0)
```

---

mhn_mean *Mean of the Modified Half-Normal Distribution*

Description

Computes $E(X)$ for $X \sim \mathrm{MHN}(\alpha, \beta, \gamma)$.

Usage

```
mhn_mean(alpha, beta, gamma)
```

Arguments

alpha	Shape parameter ($\alpha > 0$).
beta	Scale parameter ($\beta > 0$).
gamma	Location parameter ($\gamma \in R$).

Details

The mean is computed as a ratio of Fox-Wright Psi functions:

$$E(X) = \frac{\Psi[(\alpha+1)/2,\ \gamma/\sqrt{\beta}]}{\sqrt{\beta}\,\Psi[\alpha/2,\ \gamma/\sqrt{\beta}]}$$

Value

A numeric scalar.

References

Sun, J., Kong, M., & Pal, S. (2023). The Modified-Half-Normal distribution: Properties and an efficient sampling scheme. *Communications in Statistics - Theory and Methods*, 52(5), 1507–1536. (Lemma 2a)

See Also

mhn_var, dmhn

Examples

```
mhn_mean(alpha = 2, beta = 1, gamma = 0)
```

---

mhn_mode *Mode of the Modified Half-Normal Distribution*

Description

Computes the mode (most probable value) of the MHN distribution.

Usage

```
mhn_mode(alpha, beta, gamma)
```

Arguments

alpha	Shape parameter ($\alpha > 0$).
beta	Scale parameter ($\beta > 0$).
gamma	Location parameter ($\gamma \in R$).

Details

The mode depends on α:

$\alpha > 1$ $(\gamma + \sqrt{\gamma^2 + 8\beta(\alpha - 1)})/(4\beta)$ (Sun et al., 2023, Lemma 3b).

$\alpha = 1$ $\max(0, \gamma/(2\beta))$, obtained as the mode of the truncated normal $\mathrm{TN}(\gamma/(2\beta), 1/\sqrt{2\beta}, 0, \infty)$ that the MHN reduces to in this case (Sun et al., 2023, Lemma 6b).

$0 < \alpha < 1$ An interior mode exists only when $\gamma > 0$ and $\alpha \geq 1 - \gamma^2/(8\beta)$ (Sun et al., 2023, Lemma 3c); otherwise the density is monotonically decreasing (Sun et al., 2023, Lemma 3d) and NA is returned.

Value

A numeric scalar. Returns NA when no interior mode exists (density is monotonically decreasing on $(0, \infty)$).

References

Sun, J., Kong, M., & Pal, S. (2023). The Modified-Half-Normal distribution: Properties and an efficient sampling scheme. *Communications in Statistics - Theory and Methods*, 52(5), 1507–1536. (Lemma 3b–d, Lemma 6b)

See Also

dmhn, mhn_mean

Examples

```
mhn_mode(alpha = 2, beta = 1, gamma = 1)
mhn_mode(alpha = 1, beta = 1, gamma = 2)
mhn_mode(alpha = 0.5, beta = 1, gamma = -1)  # NA
```

---

mhn_skewness *Skewness of the Modified Half-Normal Distribution*

Description

Computes the skewness $\gamma_1 = E[(X - \mu)^3]/\sigma^3$ for $X \sim \mathrm{MHN}(\alpha, \beta, \gamma)$.

Usage

```
mhn_skewness(alpha, beta, gamma)
```

Arguments

alpha	Shape parameter ($\alpha > 0$).
beta	Scale parameter ($\beta > 0$).
gamma	Location parameter ($\gamma \in R$).

Details

Uses the moment recurrence (Sun et al., 2023, Lemma 2b) to compute raw moments up to third order, then converts to central moments.

Value

A numeric scalar.

References

Sun, J., Kong, M., & Pal, S. (2023). The Modified-Half-Normal distribution: Properties and an efficient sampling scheme. *Communications in Statistics - Theory and Methods*, 52(5), 1507–1536. (Lemma 2b)

See Also

`mhn_kurtosis`, `mhn_mean`

Examples

```
mhn_skewness(alpha = 2, beta = 1, gamma = 0)
```

`mhn_var` *Variance of the Modified Half-Normal Distribution*

Description

Computes $\mathrm{Var}(X)$ for $X \sim \mathrm{MHN}(\alpha, \beta, \gamma)$.

Usage

```
mhn_var(alpha, beta, gamma)
```

Arguments

`alpha`	Shape parameter ($\alpha > 0$).
`beta`	Scale parameter ($\beta > 0$).
`gamma`	Location parameter ($\gamma \in R$).

Details

Uses the formula (Sun et al., 2023, Lemma 2c):

$$\mathrm{Var}(X) = \frac{\alpha}{2\beta} + E(X)\left(\frac{\gamma}{2\beta} - E(X)\right)$$

For $\alpha \geq 1$, the variance satisfies $\mathrm{Var}(X) \leq 1/(2\beta)$ (Sun et al., 2023, Lemma 4c).

Value

A numeric scalar.

References

Sun, J., Kong, M., & Pal, S. (2023). The Modified-Half-Normal distribution: Properties and an efficient sampling scheme. *Communications in Statistics - Theory and Methods*, 52(5), 1507–1536. (Lemma 2c)

See Also

mhn_mean, dmhn

Examples

```
mhn_var(alpha = 2, beta = 1, gamma = 0)
```

---

pmhn *Distribution Function of the Modified Half-Normal Distribution*

Description

Computes the cumulative distribution function (CDF) of the Modified Half-Normal (MHN) distribution with parameters `alpha`, `beta`, and `gamma`.

Usage

```
pmhn(q, alpha = 1, beta = 1, gamma = 0, lower.tail = TRUE, log.p = FALSE)
```

Arguments

`q`	Numeric vector of quantiles.
`alpha`	Shape parameter ($\alpha > 0$). Scalar or numeric vector. Default: 1.
`beta`	Scale parameter ($\beta > 0$). Scalar or numeric vector. Default: 1.
`gamma`	Location parameter ($\gamma \in R$). Scalar or numeric vector. Default: 0.
`lower.tail`	Logical; if `TRUE` (default), probabilities are $P(X \leq q)$, otherwise $P(X > q)$.
`log.p`	Logical; if `TRUE`, probabilities are returned on the log scale. Default: `FALSE`.

Details

The CDF is computed via the series representation

$$F(x \mid \alpha, \beta, \gamma) = \frac{1}{\Psi[\alpha/2, \gamma/\sqrt{\beta}]} \sum_{i=0}^{\infty} \frac{z^i}{i!} \Gamma(s_i) \, P(s_i, \beta x^2)$$

where $z = \gamma/\sqrt{\beta}$, $s_i = (\alpha + i)/2$, and $P(s, y)$ is the regularized lower incomplete gamma function (Sun et al., 2023, Lemma 1b; equivalent to the paper's form via the identity $\Gamma(s)\, P(s, y) = \gamma(s, y)$, where $\gamma(s, y)$ is the lower incomplete gamma function used in the paper). The infinite sum is truncated at the constructive bound $K = \max\{K_1, K_2\}$ from Sun et al. (2023), Supplementary Lemma 10(d), which makes the truncation residual bounded by the user's tolerance divided by Ψ. When double-precision cancellation in the alternating-sign accumulator for $\gamma < 0$ would exceed that tolerance, the series is replaced by a Gauss-Kronrod (or tanh-sinh for $\alpha < 1$) numerical integration of the density on $[0, q]$.

Special cases are detected and dispatched to standard R primitives:

- $\gamma = 0$: `pgamma(q^2, alpha/2, scale = 1/beta)`
- $\alpha = 1$: truncated-normal CDF via `pnorm`

When any of `alpha`, `beta`, `gamma` is a vector, the CDF is evaluated element-wise. The Fox-Wright Ψ normalizing constant is recomputed only when consecutive elements present a different (α, β, γ) triple, so passing grouped parameters is significantly faster than calling `pmhn` inside an R loop.

Value

A numeric vector. The output length equals `max(length(q), length(alpha), length(beta), length(gamma))`; each input is recycled to that length following standard R recycling rules. For `q <= 0` the CDF is 0; for `q = Inf` it is 1.

References

Sun, J., Kong, M., & Pal, S. (2023). The Modified-Half-Normal distribution: Properties and an efficient sampling scheme. *Communications in Statistics - Theory and Methods*, 52(5), 1507–1536.

See Also

`dmhn`, `qmhn`, `rmhn`

Examples

```
# Basic evaluation
pmhn(c(0.5, 1, 1.5), alpha = 2, beta = 1, gamma = 1)

# Tail / log forms
pmhn(2, alpha = 2, beta = 1, gamma = 1, lower.tail = FALSE)
pmhn(2, alpha = 2, beta = 1, gamma = 1, log.p = TRUE)

# Special case: gamma = 0 reduces to sqrt-Gamma
all.equal(pmhn(1.5, alpha = 2, beta = 1, gamma = 0),
          pgamma(1.5^2, shape = 1, rate = 1))
```

---

qmhn *Quantile Function of the Modified Half-Normal Distribution*

Description

Computes the quantile (inverse cumulative) function of the Modified Half-Normal (MHN) distribution with parameters `alpha`, `beta`, and `gamma`.

Usage

```
qmhn(p, alpha = 1, beta = 1, gamma = 0, lower.tail = TRUE, log.p = FALSE)
```

Arguments

`p`	Numeric vector of probabilities.
`alpha`	Shape parameter ($\alpha > 0$). Scalar or numeric vector. Default: 1.
`beta`	Scale parameter ($\beta > 0$). Scalar or numeric vector. Default: 1.
`gamma`	Location parameter ($\gamma \in R$). Scalar or numeric vector. Default: 0.
`lower.tail`	Logical; if `TRUE` (default), probabilities are $P(X \leq q)$, otherwise $P(X > q)$.
`log.p`	Logical; if `TRUE`, probabilities are provided on the log scale. Default: `FALSE`.

Details

For the general case, $q = F^{-1}(p)$ is obtained by a TOMS 748 root-finder applied to the series CDF (Sun et al., 2023, Lemma 1b). The initial bracket is $[\sqrt{\epsilon},\ E(X) + 8\sqrt{\mathrm{Var}(X)}]$ and is doubled on the right (up to 30 times) until it brackets the target probability.

Special cases are detected and dispatched to standard R primitives:

- $\gamma = 0$: `sqrt(qgamma(p, alpha/2, scale = 1/beta))`
- $\alpha = 1$: truncated-normal inverse via `qnorm`

When any of `alpha`, `beta`, `gamma` is a vector, the quantile is evaluated element-wise. The Fox-Wright Ψ normalizing constant and moments $E(X)$, $\mathrm{Var}(X)$ (used to size the root-finder bracket) are recomputed only when consecutive elements present a different (α, β, γ) triple.

Value

A numeric vector. The output length equals `max(length(p), length(alpha), length(beta), length(gamma))`; each input is recycled to that length following standard R recycling rules. `qmhn(0) = 0` and `qmhn(1) = Inf`. Probabilities outside $[0, 1]$ yield `NaN`.

References

Sun, J., Kong, M., & Pal, S. (2023). The Modified-Half-Normal distribution: Properties and an efficient sampling scheme. *Communications in Statistics - Theory and Methods*, 52(5), 1507–1536.

See Also

`dmhn`, `pmhn`, `rmhn`

Examples

```
# Basic evaluation
qmhn(c(0.1, 0.5, 0.9), alpha = 2, beta = 1, gamma = 1)

# Round-trip: F(F^-1(p)) ~ p
p <- c(0.05, 0.25, 0.5, 0.75, 0.95)
all.equal(pmhn(qmhn(p, alpha = 2, beta = 1, gamma = 1),
               alpha = 2, beta = 1, gamma = 1),
          p, tolerance = 1e-6)
```

```
# Tail / log forms
qmhn(0.95, alpha = 2, beta = 1, gamma = 1, lower.tail = FALSE)
qmhn(log(0.05), alpha = 2, beta = 1, gamma = 1, log.p = TRUE)
```

---

rmhn *Random Generation from the Modified Half-Normal Distribution*

Description

Draws random variates from the Modified Half-Normal (MHN) distribution with parameters `alpha`, `beta`, and `gamma`.

Usage

```
rmhn(n, alpha = 1, beta = 1, gamma = 0, method = c("auto", "rtdr", "sun"))
```

Arguments

`n`	Non-negative integer giving the number of variates to draw. `n = 0` returns `numeric(0)`.
`alpha`	Shape parameter ($\alpha > 0$). Scalar or numeric vector. Default: 1.
`beta`	Scale parameter ($\beta > 0$). Scalar or numeric vector. Default: 1.
`gamma`	Location parameter ($\gamma \in R$). Scalar or numeric vector. Default: 0.
`method`	Sampling algorithm. One of `"auto"` (default), `"rtdr"`, or `"sun"`. See Details.

Details

The MHN density is

$$f(x \mid \alpha, \beta, \gamma) = \frac{2\beta^{\alpha/2} x^{\alpha-1} \exp(-\beta x^2 + \gamma x)}{\Psi[\alpha/2, \gamma/\sqrt{\beta}]} \quad (x > 0)$$

where $\Psi[a, z]$ is the Fox-Wright Psi function. `rmhn` does not evaluate Ψ; the rejection-sampling kernels cancel it out.

The default parameters `alpha = 1, beta = 1, gamma = 0` correspond to the half-normal distribution $\mathrm{HN}(1/\sqrt{2})$.

The `method` argument selects the rejection sampler:

- `"auto"`: Special-case shortcuts when applicable ($\gamma \approx 0$ -> sqrt-Gamma, $\alpha \approx 1$ -> truncated normal). Otherwise dispatches to RTDR (Gao & Wang, 2025).
- `"rtdr"`: Force the Relaxed Transformed Density Rejection method of Gao & Wang (2025). The acceptance probability is bounded below by $1/e \approx 0.368$ uniformly over the parameter space. Note: Gao & Wang (2025) use the parameterization (λ, α, β) with density proportional to $x^{\lambda-1} \exp(-\alpha x^2 - \beta x)$; the mapping to the Sun et al. parameterization used here is $\lambda \leftrightarrow \alpha$, $\alpha \leftrightarrow \beta$, $\beta \leftrightarrow -\gamma$ (sign flip on the linear term).

- "sun": Force the Sun et al. (2023) algorithms. Algorithm 1 is used when $\gamma > 0$ and $\alpha > 1$; Algorithm 3 is used when $\gamma \leq 0$. The combination $\alpha < 1$ with $\gamma > 0$ is unsupported and triggers an error.

Vector parameters are recycled to length n following standard R rules. Trailing parameter elements beyond index n - 1 are silently ignored, matching the convention of rnorm.

Internally the setup state of the chosen sampler is reused as long as consecutive (α, β, γ) triples are equal, so passing parameters grouped by triple is faster than calling rmhn inside an R loop.

Value

A numeric vector of length n. If any of alpha, beta, gamma (after recycling to length n) is NA or non-finite (Inf, -Inf, NaN), the corresponding output element is NA.

References

Sun, J., Kong, M., & Pal, S. (2023). The Modified-Half-Normal distribution: Properties and an efficient sampling scheme. *Communications in Statistics - Theory and Methods*, 52(5), 1507–1536.

Gao, F. & Wang, H.-B. (2025). Generating modified-half-normal random variates by a relaxed transformed density rejection method. *Communications in Statistics - Simulation and Computation*.

Robert, C. P. (1995). Simulation of truncated normal variables. *Statistics and Computing*, 5(2), 121–125.

See Also

dmhn, mhn_mean, mhn_var

Examples

```
set.seed(1)
rmhn(10, alpha = 2, beta = 1, gamma = 0.5)

# Vector parameters are recycled to length n.
set.seed(1)
rmhn(5, alpha = c(1, 2, 3, 4, 5))
```

Index