

Package ‘bqmm’

July 10, 2026

Title Bayesian Multilevel Quantile Regression

Version 0.1.0

Description Fits Bayesian mixed-effects (multilevel) quantile regression models using the asymmetric Laplace working likelihood and Stan. Supports an 'lme4'-style formula interface with nested and crossed random effects, fitting one or several quantiles, post-hoc non-crossing rearrangement of fitted quantiles, and the Yang, Wang and He (2016) posterior-variance correction for valid frequentist inference from the (misspecified) asymmetric Laplace posterior.

License MIT + file LICENSE

Encoding UTF-8

RoxygenNote 7.3.3

Biarch true

Depends R (>= 4.0.0)

Imports methods, stats, graphics, grDevices, lme4, Matrix, Rcpp (>= 0.12.0), RcppParallel (>= 5.0.1), rstan (>= 2.26.0), rstantools (>= 2.4.0), posterior

LinkingTo BH (>= 1.66.0), Rcpp (>= 0.12.0), RcppEigen (>= 0.3.3.3.0), RcppParallel (>= 5.0.1), rstan (>= 2.26.0), StanHeaders (>= 2.26.0)

Suggests bayesplot, ggplot2, lqmm, MASS, nlme, quantreg, testthat (>= 3.0.0), knitr, rmarkdown

SystemRequirements GNU make

Config/testthat/edition 3

Config/Needs/website pkgdown

VignetteBuilder knitr

URL <https://github.com/kvenkita/bqmm>, <https://kvenkita.github.io/bqmm/>

BugReports <https://github.com/kvenkita/bqmm/issues>

NeedsCompilation yes

Author Kailas Venkitasubramanian [aut, cre, cph]

Maintainer Kailas Venkitasubramanian <kailasv@gmail.com>

Repository CRAN

Date/Publication 2026-07-10 20:00:02 UTC

Contents

ald	2
as.matrix.bqmm	3
as_draws.bqmm	4
bqmm	4
bqmm_prior	6
coef.bqmm	7
coef.bqmm_multi	7
confint.bqmm	8
fitted.bqmm	9
fixef.bqmm	9
log_lik.bqmm	10
nobs.bqmm	11
plot.bqmm	11
plot.bqmm_multi	12
posterior_epred.bqmm	13
posterior_predict.bqmm	13
predict.bqmm	14
ranef.bqmm	15
rearrange_quantiles	15
residuals.bqmm	16
summary.bqmm	17
summary.bqmm_multi	18
VarCorr.bqmm	18
vcov.bqmm	19
Index	20

ald

The asymmetric Laplace family for quantile regression

Description

A lightweight family object describing the asymmetric Laplace distribution (ALD) working likelihood used by `bqmm()`. It mirrors the role of a `stats::family` object but is intentionally minimal in this release.

Usage

```
ald(link = "identity")
```

Arguments

link Name of the link for the location (quantile) parameter. Only "identity" is supported in v0.1.

Value

An object of class "bqmm_family".

Examples

```
ald()
```

as.matrix.bqmm *Coerce a bqmm fit to a matrix of posterior draws*

Description

Coerce a bqmm fit to a matrix of posterior draws

Usage

```
## S3 method for class 'bqmm'
as.matrix(x, ...)
```

Arguments

x A bqmm fit.
... Unused.

Value

A matrix of posterior draws (draws in rows, parameters in columns), using the raw Stan parameter names.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)
dim(as.matrix(fit))
```

as_draws.bqmm	<i>Convert a bqmm fit to a posterior draws object</i>
---------------	---

Description

Convert a bqmm fit to a posterior draws object

Usage

```
## S3 method for class 'bqmm'
as_draws(x, ...)
```

Arguments

x	A bqmm fit.
...	Unused.

Value

A draws_array (from the posterior package) with tidy variable names: b_<name> for fixed effects, sd_<component> for random-effect SDs, and sigma.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)
as_draws(fit)
```

bqmm	<i>Bayesian multilevel quantile regression</i>
------	--

Description

Fits a Bayesian mixed-effects quantile regression model using the asymmetric Laplace working likelihood and Stan. The interface follows lme4: random effects are written inline in the formula, e.g. $y \sim x + (1 + x \mid \text{group})$, and nested or crossed grouping factors are both supported.

Usage

```

bqmm(
  formula,
  data,
  tau = 0.5,
  family = ald(),
  prior = NULL,
  cov = c("diagonal", "unstructured"),
  adjust = TRUE,
  prior_only = FALSE,
  chains = 4,
  iter = 2000,
  warmup = floor(iter/2),
  cores = getOption("mc.cores", 1L),
  seed = NULL,
  control = list(adapt_delta = 0.95),
  ...
)

```

Arguments

formula	An lme4-style model formula.
data	A data frame containing the variables in formula.
tau	Quantile level(s) in (0, 1). Scalar or vector.
family	A <code>bqmm_family</code> object; currently only <code>ald()</code> .
prior	A <code>bqmm_prior()</code> object, or NULL for data-scaled defaults.
cov	Random-effect covariance structure. "diagonal" (default) models independent random effects and supports any number of nested or crossed terms. "unstructured" adds an LKJ-correlated covariance but currently requires exactly one random-effects term (e.g. $y \sim x + (1 + x g)$).
adjust	Logical; compute the Yang-Wang-He (2016) variance correction so that <code>vcov(fit, adjusted = TRUE)</code> returns valid fixed-effect uncertainty. Default TRUE.
prior_only	Logical; sample from the prior predictive distribution.
chains, iter, warmup, cores, seed	Passed to <code>rstan::sampling()</code> .
control	A list of sampler control parameters (e.g. <code>adapt_delta</code>). Defaults raise <code>adapt_delta</code> to 0.95 because ALD posteriors are sharp.
...	Additional arguments forwarded to <code>rstan::sampling()</code> .

Details

One or several quantiles may be requested through `tau`. A scalar returns a single `bqmm` fit; a vector fits each quantile independently and returns a `bqmm_multi` container.

Value

A `bqmm` object (single `tau`) or a `bqmm_multi` object (vector `tau`).

Examples

```
# A minimal fit; raise chains/iter for real analyses.
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)
summary(fit)
```

bqmm_prior

*Priors for a Bayesian quantile mixed model***Description**

Builds the list of prior hyperparameters passed to Stan. Defaults are weakly informative and scaled to the data (see `bqmm_default_priors()`); any element supplied here overrides the default.

Usage

```
bqmm_prior(
  beta_mean = 0,
  beta_sd = NULL,
  sigma_scale = NULL,
  re_scale = NULL,
  lkj = 2
)
```

Arguments

beta_mean	Numeric scalar or vector: prior mean(s) for the fixed-effect coefficients. Recycled to the number of columns of the design matrix. Default 0.
beta_sd	Positive scalar or vector: prior SD(s) for the fixed-effect coefficients. NULL (default) uses a data-scaled value.
sigma_scale	Positive scalar: half-normal scale for the ALD scale sigma. NULL (default) uses a data-scaled value.
re_scale	Positive scalar: half-normal scale for the random-effect standard deviations. NULL (default) uses a data-scaled value.
lkj	Positive scalar: LKJ shape parameter for the random-effect correlation matrix (used only when cov = "unstructured"). 2 favours weak correlations; 1 is uniform over correlation matrices.

Value

An object of class "bqmm_prior".

Examples

```
bqmm_prior(beta_sd = 5)
```

coef.bqmm	<i>Extract model coefficients</i>
-----------	-----------------------------------

Description

Alias for `fixef()`; returns the posterior-median fixed effects.

Usage

```
## S3 method for class 'bqmm'
coef(object, ...)
```

Arguments

object	A bqmm fit.
...	Unused.

Value

A named numeric vector of posterior-median fixed-effect coefficients.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)
coef(fit)
```

coef.bqmm_multi	<i>Coefficient-versus-tau matrix for a bqmm_multi fit</i>
-----------------	---

Description

Coefficient-versus-tau matrix for a bqmm_multi fit

Usage

```
## S3 method for class 'bqmm_multi'
coef(object, ...)
```

Arguments

object	A bqmm_multi fit.
...	Unused.

Value

A tau-by-coefficient matrix of posterior-median fixed effects, with one row per quantile.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = c(0.25, 0.75), chains = 1, iter = 250,
            refresh = 0, seed = 1)
coef(fit)
```

confint.bqmm

Confidence (credible) intervals for the fixed effects

Description

Wald-type intervals built from the posterior-median estimates and the (optionally misspecification-corrected) fixed-effect covariance.

Usage

```
## S3 method for class 'bqmm'
confint(
  object,
  parm,
  level = 0.95,
  adjusted = TRUE,
  method = c("ywh", "ij"),
  cluster = TRUE,
  ...
)
```

Arguments

object	A bqmm fit.
parm	Optional subset of coefficients (names or indices) to return.
level	Interval coverage (default 0.95).
adjusted	Logical; if TRUE (default) use the corrected covariance from <code>vcov.bqmm()</code> , otherwise the naive posterior covariance.
method	Correction to use when adjusted = TRUE; see <code>vcov.bqmm()</code> .
cluster	Logical; use the cluster-robust form (default TRUE).
...	Unused.

Value

A matrix with one row per coefficient and lower/upper interval columns.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)
confint(fit)
```

fitted.bqmm	<i>Linear predictor (conditional tau-quantile) at the posterior median</i>
-------------	--

Description

Linear predictor (conditional tau-quantile) at the posterior median

Usage

```
## S3 method for class 'bqmm'
fitted(object, ...)
```

Arguments

object	A bqmm fit.
...	Unused.

Value

Numeric vector of fitted conditional quantiles.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)
head(fitted(fit))
```

fixef.bqmm	<i>Posterior-median fixed effects</i>
------------	---------------------------------------

Description

Posterior-median fixed effects

Usage

```
## S3 method for class 'bqmm'
fixef(object, ...)
```

Arguments

object	A bqmm fit.
...	Unused.

Value

A named numeric vector of posterior-median fixed-effect coefficients.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)
fixef(fit)
```

log_lik.bqmm

Pointwise log-likelihood draws

Description

Pointwise log-likelihood draws

Usage

```
## S3 method for class 'bqmm'
log_lik(object, ...)
```

Arguments

object	A bqmm fit.
...	Unused.

Value

An S x N matrix of pointwise log-likelihood values, suitable for use with the loo package.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)
dim(log_lik(fit))
```

nobs.bqmm	<i>Number of observations used in the fit</i>
-----------	---

Description

Number of observations used in the fit

Usage

```
## S3 method for class 'bqmm'
nobs(object, ...)
```

Arguments

object	A bqmm fit.
...	Unused.

Value

An integer, the number of observations.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)
nobs(fit)
```

plot.bqmm	<i>Plot a bqmm fit</i>
-----------	------------------------

Description

Default plot shows fixed-effect posterior intervals. With bayesplot installed, richer MCMC plots are available via [as_draws\(\)](#).

Usage

```
## S3 method for class 'bqmm'
plot(x, ...)
```

Arguments

x	A bqmm fit.
...	Unused.

Value

Invisibly, x.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)
plot(fit)
```

plot.bqmm_multi	<i>Plot coefficient-versus-tau paths for a bqmm_multi fit</i>
-----------------	---

Description

Plot coefficient-versus-tau paths for a bqmm_multi fit

Usage

```
## S3 method for class 'bqmm_multi'
plot(x, ...)
```

Arguments

x	A bqmm_multi fit.
...	Unused.

Value

Invisibly, x.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = c(0.25, 0.75), chains = 1, iter = 250,
            refresh = 0, seed = 1)
plot(fit)
```

posterior_epred.bqmm *Draws of the expected response (conditional tau-quantile)*

Description

Draws of the expected response (conditional tau-quantile)

Usage

```
## S3 method for class 'bqmm'  
posterior_epred(object, ...)
```

Arguments

object	A bqmm fit.
...	Unused.

Value

An $S \times N$ matrix of posterior draws of the linear predictor, where S is the number of posterior draws and N the number of observations.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,  
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)  
dim(posterior_epred(fit))
```

posterior_predict.bqmm

Draws from the posterior predictive distribution

Description

Draws from the posterior predictive distribution

Usage

```
## S3 method for class 'bqmm'  
posterior_predict(object, ...)
```

Arguments

object	A bqmm fit.
...	Unused.

Value

An $S \times N$ matrix of posterior predictive draws of the response, where S is the number of posterior draws and N the number of observations.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)
dim(posterior_predict(fit))
```

predict.bqmm	<i>Predictions from a bqmm fit</i>
--------------	------------------------------------

Description

Predictions from a bqmm fit

Usage

```
## S3 method for class 'bqmm'
predict(
  object,
  newdata = NULL,
  re.form = NULL,
  noncrossing = c("none", "rearrange"),
  ...
)
```

Arguments

object	A bqmm fit.
newdata	Optional data frame; if omitted, training data are used.
re.form	NULL includes random effects (training data only); NA gives population-level predictions.
noncrossing	One of "none" or "rearrange". Rearrangement only has an effect for bqmm_multi objects (multiple quantiles).
...	Unused.

Value

Numeric vector of predicted conditional quantiles.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)
head(predict(fit, re.form = NA))
```

ranef.bqmm	<i>Posterior-median random effects</i>
------------	--

Description

Posterior-median random effects

Usage

```
## S3 method for class 'bqmm'
ranef(object, ...)
```

Arguments

object	A bqmm fit.
...	Unused.

Value

A numeric vector of posterior-median random effects aligned with the columns of the random-effects design matrix Z, or NULL if the model has no random effects.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)
ranef(fit)
```

rearrange_quantiles	<i>Rearrange fitted quantiles to remove crossing</i>
---------------------	--

Description

Post-hoc monotonisation of estimated quantile curves (Chernozhukov, Fernandez-Val and Galichon, 2010): at any covariate point, sorting the fitted values across quantile levels into increasing order yields a valid, non-crossing set of quantiles and never increases estimation error. This is the v0.1 non-crossing strategy; joint constrained estimation is deferred.

Usage

```
rearrange_quantiles(preds)
```

Arguments

`preds` A numeric matrix of fitted quantiles with one column per quantile level, ordered by increasing tau (rows = observations).

Value

A matrix of the same shape with each row sorted increasingly.

Examples

```
m <- rbind(c(1, 0.5, 2), c(0, 1, 0.8)) # some crossings
rearrange_quantiles(m)
```

residuals.bqmm	<i>Residuals from a bqmm fit</i>
----------------	----------------------------------

Description

Residuals from a bqmm fit

Usage

```
## S3 method for class 'bqmm'
residuals(object, ...)
```

Arguments

`object` A bqmm fit.
`...` Unused.

Value

Numeric vector of response-minus-fitted residuals.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)
head(residuals(fit))
```

summary.bqmm	<i>Summarize a bqmm fit</i>
--------------	-----------------------------

Description

Produces a fixed-effect coefficient table (estimate, standard error and interval) together with random-effect standard deviations.

Usage

```
## S3 method for class 'bqmm'
summary(
  object,
  level = 0.95,
  adjusted = TRUE,
  method = c("ywh", "ij"),
  cluster = TRUE,
  ...
)
```

Arguments

object	A bqmm fit.
level	Interval coverage (default 0.95).
adjusted	Logical; if TRUE (default) use the corrected covariance from vcov.bqmm() for the standard errors and intervals.
method	Correction to use when adjusted = TRUE; see vcov.bqmm() .
cluster	Logical; use the cluster-robust form (default TRUE).
...	Unused.

Value

An object of class `summary.bqmm`.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)
summary(fit)
```

summary.bqmm_multi *Summarize a bqmm_multi fit*

Description

Summarize a bqmm_multi fit

Usage

```
## S3 method for class 'bqmm_multi'
summary(object, ...)
```

Arguments

object A bqmm_multi fit.
 ... Passed to the per-quantile summary() method for each fit.

Value

A list of summary.bqmm objects, one per quantile.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = c(0.25, 0.75), chains = 1, iter = 250,
            refresh = 0, seed = 1)
summary(fit)
```

VarCorr.bqmm *Random-effect standard deviations and correlations*

Description

Random-effect standard deviations and correlations

Usage

```
## S3 method for class 'bqmm'
VarCorr(x, sigma = 1, ...)
```

Arguments

x A bqmm fit.
 sigma Ignored; present for compatibility with the generic.
 ... Unused.

Value

A named numeric vector of posterior-median random-effect standard deviations (with a posterior-median correlation matrix attached as the "correlation" attribute for unstructured models), or NULL if the model has no random effects.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)
VarCorr(fit)
```

vcov.bqmm

Variance-covariance of the fixed effects

Description

Variance-covariance of the fixed effects

Usage

```
## S3 method for class 'bqmm'
vcov(object, adjusted = TRUE, method = c("ywh", "ij"), cluster = TRUE, ...)
```

Arguments

object	A bqmm fit.
adjusted	Logical; if TRUE (default) return a misspecification- corrected covariance (chosen by method), otherwise the naive posterior covariance.
method	Correction to use when adjusted = TRUE: "ywh" (default) is the Yang-Wang-He posterior-covariance sandwich (<code>compute_ywh_multiplicative()</code>); "ij" is the Infinitesimal Jackknife (<code>ij_vcov()</code>). Both are cluster-robust by default for a mixed model.
cluster	Logical; use the cluster-robust form (default TRUE).
...	Unused.

Value

A $K \times K$ covariance matrix for the fixed effects.

Examples

```
fit <- bqmm(distance ~ age + (1 | Subject), data = nlme::Orthodont,
            tau = 0.5, chains = 1, iter = 300, refresh = 0, seed = 1)
vcov(fit)
```

Index

ald, 2
ald(), 5
as.matrix.bqmm, 3
as_draws(), 11
as_draws.bqmm, 4

bqmm, 4
bqmm(), 2
bqmm_default_priors(), 6
bqmm_prior, 6
bqmm_prior(), 5

coef.bqmm, 7
coef.bqmm_multi, 7
compute_ywh_multiplicative(), 19
confint.bqmm, 8

fitted.bqmm, 9
fixef.bqmm, 9

ij_vcov(), 19

log_lik.bqmm, 10

nobs.bqmm, 11

plot.bqmm, 11
plot.bqmm_multi, 12
posterior_epred.bqmm, 13
posterior_predict.bqmm, 13
predict.bqmm, 14

ranef.bqmm, 15
rearrange_quantiles, 15
residuals.bqmm, 16
rstan::sampling(), 5

stats::family, 2
summary.bqmm, 17
summary.bqmm_multi, 18

VarCorr.bqmm, 18
vcov.bqmm, 19
vcov.bqmm(), 8, 17