

Package ‘ordbetareg’

March 15, 2023

Type Package

Title Ordered Beta Regression Models with 'brms'

Version 0.7.1

Description Implements ordered beta regression models, which are for modeling continuous variables with upper and lower bounds, such as survey sliders, dose-response relationships and indexes. For more information, see Kubinec (2022) <[doi:10.31235/osf.io/2sx6y](https://doi.org/10.31235/osf.io/2sx6y)>. The package is a front-end to the R package 'brms', which facilitates a range of regression specifications, including hierarchical, dynamic and multivariate modeling.

BugReports https://github.com/saudiwin/ordbetareg_pack/issues

License MIT + file LICENSE

Encoding UTF-8

LazyData true

LazyDataCompression xz

RoxygenNote 7.2.3

Depends R (>= 3.5), brms (>= 2.18.0), stats

Imports transformr, dplyr, ggplot2 (>= 3.4.0), ganimate, tidyr

Suggests rmarkdown, knitr, modelsummary, marginaleffects (>= 0.10.0), haven, stringr, Hmisc, collapse, ggthemes, glmmTMB, mice, bayestestR, kableExtra

VignetteBuilder knitr

NeedsCompilation no

Author Robert Kubinec [aut, cre] (<<https://orcid.org/0000-0001-6655-4119>>)

Maintainer Robert Kubinec <rmk7@nyu.edu>

Repository CRAN

Date/Publication 2023-03-15 18:10:02 UTC

R topics documented:

dordbeta	2
fit_imputed	3
fit_multivariate	3
normalize	4
ordbetareg	5
ord_fit_mean	8
ord_fit_phi	9
pew	9
pp_check_ordbeta	10
rordbeta	11
sim_data	12
sim_ordbeta	12
Index	15

dordbeta	<i>Probability Density Function for the Ordered Beta Distribution</i>
----------	-----------------------------------------------------------------------

Description

This function will return the density of given variates of the ordered beta distribution conditional on values for the mean (μ), dispersion (ϕ) and cutpoints governing the ratio of degenerate (discrete) to continuous responses.

Usage

```
dordbeta(x = 0.9, mu = 0.5, phi = 1, cutpoints = c(-1, 1), log = FALSE)
```

Arguments

x	Variates of the ordered beta distribution (should be in the [0,1] interval).
mu	Value of the mean of the distribution. Should be in the $\setminus(0,1\setminus)$ interval (cannot be strictly equal to 0 or 1). If length is greater than 1, should be of length x.
phi	Value of the dispersion parameter. Should be strictly greater than 0. If length is greater than 1, should be of length x.
cutpoints	A vector of two numeric values for the cutpoints. Second value should
log	where to return the log density be strictly greater than the first value.

Value

Returns a vector of length x of the density of the ordered beta distribution conditional on μ and ϕ .

Examples

```

# examine density (likelihood) of different possible values
# given fixed values for ordered beta parameters

x <- seq(0, 1, by=0.01)

x_dens <- dordbeta(x, mu = 0.3, phi=2, cutpoints=c(-2, 2))

# Most likely value for x is approx 1
# Note discontinuity in density function between continuous/discrete values
# density function is a combined PMF/PDF, so not a real PDF
# can though be used for MLE

plot(x_dens, x)

# discrete values should be compared to each other:
# prob of discrete 0 > prob of discrete 1

x_dens[x==0] > x_dens[x==1]

```

fit_imputed

Fitted Ordered Beta Regression Model (Imputed Datasets)

Description

A fitted ordered beta regression model on multiple imputed datasets generated by the package mice.

Usage

```
fit_imputed
```

Format

an ordbetareg object

fit_multivariate

Fitted Ordered Beta Regression Model (Multivariate regression)

Description

A fitted ordered beta regression model with two responses, one an ordered beta regression and the other a Gaussian/Normal outcome. Useful for examining mediation analysis.

Usage

```
fit_multivariate
```

Format

an ordbetareg object

normalize	<i>Normalize Outcome/Response to $\backslash[0,1\backslash]$ Interval</i>
-----------	--------------------------------------------------------------------------------------

Description

This function takes a continuous (double) column of data and converts it to have 0 as the lower bound and 1 as the upper bound.

Usage

```
normalize(outcome, true_bounds = NULL)
```

Arguments

outcome	Any non-character vector. Factors will be converted to numeric via coercion.
true_bounds	Specify this parameter with the lower and upper bound if the observed min/max of the outcome should not be used. Useful when an upper or lower bound exists but the observed data is less than/more than that bound. The normalization function will respect these bounds.

Details

Beta regression can only be done with a response that is continuous with a lower bound of 0 and an upper bound of 1. However, it is straightforward to transform any lower and upper-bounded continuous variable to the $\backslash[0,1\backslash]$ interval. This function does the transformation and saves the original bounds as attributes so that the bounds can be reverse-transformed.

Value

A numeric vector with an upper bound of 1 and a lower bound of 0. The original bounds are saved in the attributes "lower_bound" and "upper_bound".

Examples

```
# set up arbitrary upper and lower-bounded vector
outcome <- runif(1000, min=-33, max=445)

# normalize to  $\backslash[0,1\backslash]$ 

trans_outcome <- normalize(outcome=outcome)
summary(trans_outcome)

# only works with numeric vectors and factors

try(normalize(outcome=c('a','b')))
```

Description

This function allows you to estimate an ordered beta regression model via a formula syntax.

The `ordbetareg` package is essentially a wrapper around `brms` that enables the ordered beta regression model to be fit. This model has advantages over other alternatives for continuous data with upper and lower bounds, such as survey sliders, indexes, dose-response relationships, and visual analog scales (among others). The package allows for all of the many `brms` regression modeling functions to be used with the ordered beta regression distribution.

Usage

```
ordbetareg(  
  formula = NULL,  
  data = NULL,  
  true_bounds = NULL,  
  phi_reg = "none",  
  use_brm_multiple = FALSE,  
  coef_prior_mean = 0,  
  coef_prior_SD = 5,  
  intercept_prior_mean = NULL,  
  intercept_prior_SD = NULL,  
  phi_prior = 0.1,  
  dirichlet_prior = c(1, 1, 1),  
  phi_coef_prior_mean = 0,  
  phi_coef_prior_SD = 5,  
  phi_intercept_prior_mean = NULL,  
  phi_intercept_prior_SD = NULL,  
  extra_prior = NULL,  
  init = "0",  
  offset = 0,  
  ...  
)
```

Arguments

- | | |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>formula</code> | Either an R formula in the form <code>response/DV ~ var1 + var2</code> etc. <i>or</i> formula object as created/called by the <code>brms::bf</code> function. *Please avoid using 0 or Intercept in the formula definition. |
| <code>data</code> | An R data frame or tibble containing the variables in the formula |
| <code>true_bounds</code> | If the true bounds of the outcome/response don't exist in the data, pass a length 2 numeric vector of the minimum and maximum bounds to properly normalize the outcome/response |

<code>phi_reg</code>	Whether you are including a linear model predicting the dispersion parameter, phi, and/or for the response. If you are including models for both, pass option 'both'. If you only have an intercept for the outcome (i.e. a 1 in place of covariates), pass 'only'. If no model for phi, the default, pass 'none'.
<code>use_brm_multiple</code>	(T/F) Whether the model should use <code>brms::brm_multiple</code> for multiple imputation over multiple dataframes passed as a list to the <code>data</code> argument
<code>coef_prior_mean</code>	The mean of the Normal distribution prior on the regression coefficients (for predicting the mean of the response). Default is 0.
<code>coef_prior_SD</code>	The SD of the Normal distribution prior on the regression coefficients (for predicting the mean of the response). Default is 5, which makes the prior weakly informative on the logit scale.
<code>intercept_prior_mean</code>	The mean of the Normal distribution prior for the intercept. By default is NULL, which means the intercept receives the same prior as <code>coef_prior_mean</code> . To zero out the intercept, set this parameter to 0 and <code>coef_prior_SD</code> to a very small number (0.01 or smaller).
<code>intercept_prior_SD</code>	The SD of the Normal distribution prior for the intercept. By default is NULL, which means the intercept receives the same prior SD as <code>coef_prior_SD</code> .
<code>phi_prior</code>	The mean parameter of the exponential prior on phi, which determines the dispersion of the beta distribution. The default is .1, which equals a mean of 10 and is thus weakly informative on the interval (0.4, 30). If the response has very low variance (i.e. tightly) clusters around a specific value, then decreasing this prior (and increasing the expected value) may be helpful. Checking the value of phi in the output of the model command will reveal if a value of 0.1 (mean of 10) is too small.
<code>dirichlet_prior</code>	A vector of three integers corresponding to the prior parameters for the dirchlet distribution (alpha parameter) governing the location of the cutpoints between the components of the response (continuous vs. degenerate). The default is 1 which puts equal probability on degenerate versus continuous responses. Likely only needs to be changed in a repeated sampling situation to stabilize the cutpoint locations across samples.
<code>phi_coef_prior_mean</code>	The mean of the Normal distribution prior on the regression coefficients for predicting phi, the dispersion parameter. Only useful if a linear model is being fit to phi. Default is 0.
<code>phi_coef_prior_SD</code>	The SD of the Normal distribution prior on the regression coefficients for predicting phi, the dispersion parameter. Only useful if a linear model is being fit to phi. Default is 5, which makes the prior weakly informative on the exponential scale.
<code>phi_intercept_prior_mean</code>	The mean of the Normal distribution prior for the phi (dispersion) regression intercept. By default is NULL, which means the intercept receives the same

	prior as <code>phi_coef_prior_mean</code> . To zero out the intercept, set this parameter to 0 and <code>phi_coef_prior_SD</code> to a very small number (0.01 or smaller).
<code>phi_intercept_prior_SD</code>	The SD of the Normal distribution prior for the phi (dispersion) regression intercept. By default is NULL, which means the intercept receives the same prior SD as <code>phi_coef_prior_SD</code> .
<code>extra_prior</code>	An additional prior, such as a prior for a specific regression coefficient, added to the outcome regression by passing one of the brms functions <code>brms::set_prior</code> or <code>brms::prior_string</code> with appropriate values.
<code>init</code>	This parameter is used to determine starting values for the Stan sampler to begin Markov Chain Monte Carlo sampling. It is set by default at 0 because the non-linear nature of beta regression means that it is possible to begin with extreme values depending on the scale of the covariates. Setting this to 0 helps the sampler find starting values. It does, on the other hand, limit the ability to detect convergence issues with Rhat statistics. If that is a concern, such as with an experimental feature of brms, set this to "random" to get more robust starting values (just be sure to scale the covariates so they are not too large in absolute size).
<code>offset</code>	Specify a hard-coded offset for the model. Default is 0 (no offset). Also consider putting an informative prior on the intercept instead.
<code>...</code>	All other arguments passed on to the brm function

Details

This function is a wrapper around the `brms::brm` function, which is a powerful Bayesian regression modeling engine using Stan. To fully explore the options available, including dynamic and hierarchical modeling, please see the documentation for the `brm` function above. As the ordered beta regression model is currently not available in brms natively, this modeling function allows a brms model to be fit with the ordered beta regression distribution.

For more information about the model, see the paper here: <https://osf.io/preprints/socarxiv/2sx6y/>.

This function allows you to set priors on the dispersion parameter, the cutpoints, and the regression coefficients (see below for options). However, to add specific priors on individual covariates, you would need to use the `brms::set_prior` function by specifying an individual covariate (see function documentation) and passing the result of the function call to the `extra_prior` argument.

This function will also automatically normalize the outcome so that it lies in the $[0,1]$ interval, as required by beta regression. For further information, see the documentation for the `normalize` function.

To learn more about how the package works, see the vignette by using the command `browseVignettes(package='ordbetareg')`

For more info about the distribution, see this paper: <https://osf.io/preprints/socarxiv/2sx6y/>

To cite the package, please cite the following paper:

Kubinec, Robert. "Ordered Beta Regression: A Parsimonious, Well-Fitting Model for Continuous Data with Lower and Upper Bounds." **Political Analysis**. 2022.

Value

A brms object fitted with the ordered beta regression distribution.

Examples

```
# load survey data that comes with the package

library(dplyr)
data("pew")

# prepare data

model_data <- select(pew, therm,
                     education="F_EDUCCAT2_FINAL",
                     region="F_CREGION_FINAL",
                     income="F_INCOME_FINAL")

# It takes a while to fit the models. Run the code
# below if you want to load a saved fitted model from the
# package, otherwise use the model-fitting code

data("ord_fit_mean")

# fit the actual model

ord_fit_mean <- ordbetareg(formula=therm ~ education + income +
                           (1|region),
                           data=model_data,
                           cores=2, chains=2)

# access values of the coefficients

summary(ord_fit_mean)
```

ord_fit_mean	<i>Fitted Ordered Beta Regression Model</i>
--------------	---------------------------------------------

Description

A fitted ordered beta regression model to the mean of the thermometer column from the pew data.

Usage

```
ord_fit_mean
```

Format

an ordbetareg object

ord_fit_phi	<i>Fitted Ordered Beta Regression Model (Phi Regression)</i>
-------------	--------------------------------------------------------------

Description

A fitted ordered beta regression model to the dispersion parameter of the thermometer column from the pew data.

Usage

```
ord_fit_phi
```

Format

an ordbetareg object

pew	<i>Pew American Trends Panel Wave 28</i>
-----	------------------------------------------

Description

A dataset with the non-missing responses for the 28th wave of the Pew American Trends Panel survey.

Usage

```
pew
```

Format

A data frame with 140 variables and 2,538 observations.

Source

<https://www.pewresearch.org/social-trends/dataset/american-trends-panel-wave-28/>

Description

The standard `brms::pp_check` plot available via `brms` is not accurate for `ordbetareg` models because an ordered beta regression has both continuous and discrete components. This function implements a bar plot and a density plot for the continuous and discrete elements separately, and will return accurate posterior predictive plots relative to the data.

Usage

```
pp_check_ordbeta(
  model = NULL,
  type = "both",
  ndraws = 10,
  cores = NULL,
  group = NULL,
  new_theme = NULL,
  outcome_label = NULL,
  animate = FALSE,
  reverse_bounds = TRUE
)
```

Arguments

<code>model</code>	A fitted <code>ordbetareg</code> model.
<code>type</code>	Default is "both" for creating both a discrete (bar) and continuous (density) plot. Can also be "discrete" for only the bar plot for discrete values (0/1) or "continuous" for continuous values (density plot).
<code>ndraws</code>	Number of posterior draws to use to calculate estimates and show in plot. Defaults to 10.
<code>cores</code>	Number of cores to use to produce posterior predictive distribution. Defaults to NULL or 1 core.
<code>group</code>	A factor variable of the same number of rows as the data that is used to produce grouped (faceted) plots of the posterior distribution.
<code>new_theme</code>	Any additional themes to be added to <code>ggplot2</code> (default is NULL).
<code>outcome_label</code>	A character value that will replace the name of the outcome in the plot (default is the name of the response variable in the data frame).
<code>animate</code>	Whether to animate each posterior draw for continuous distributions (defaults to FALSE).
<code>reverse_bounds</code>	Whether to plot data using the original bounds in the data (i.e. not 0 and 1).

Value

If "both", prints both plots and returns a list of both plots as ggplot2 objects. Otherwise, prints and returnst the specific plot as a ggplot2 object.

Examples

```
# need a fitted ordbetareg model

data("ord_fit_mean")

out_plots <- pp_check_ordbeta(ord_fit_mean)

# view discrete bar plot

out_plots$discrete

# view continuous density plot

out_plots$continuous

# change title using ggplot2 ggtitle function

out_plots$discrete + ggplot2::ggtitle("New title")
```

rordbeta

Generate Ordered Beta Variates

Description

This function will generate ordered beta random variates given values for the mean (μ), dispersion (ϕ) and cutpoints governing the ratio of degenerate (discrete) to continuous responses.

Usage

```
rordbeta(n = 100, mu = 0.5, phi = 1, cutpoints = c(-1, 1))
```

Arguments

n	Number of variates to generate.
mu	Value of the mean of the distribution. Should be in the $(0,1)$ interval (cannot be strictly equal to 0 or 1). If length is greater than 1, should be of length n.
phi	Value of the dispersion parameter. Should be strictly greater than 0. If length is greater than 1, should be of length n.
cutpoints	A vector of two numeric values for the cutpoints. Second value should be strictly greater than the first value.

Value

A vector of length n of variates from the ordered beta distribution.

Examples

```
# generate 100 random variates with an average of 0.7
# all will be in the closed interval \[0,1\]

ordbeta_var <- rordbeta(n=100, mu=0.7, phi=2)

# Will be approx mean = 0.7 with high positive skew

summary(ordbeta_var)
```

sim_data	<i>Simulated Ordered Beta Regression Values</i>
----------	-------------------------------------------------

Description

The simulated draws used in the vignette for calculating statistical power.

Usage

```
sim_data
```

Format

A dataframe

sim_ordbeta	<i>Power Calculation via Simulation of the Ordered Beta Regression Model</i>
-------------	------------------------------------------------------------------------------

Description

This function allows you to calculate power curves (or anything else) via simulating the ordered beta regression model.

Usage

```

sim_ordbeta(
  N = 1000,
  k = 5,
  iter = 1000,
  cores = 1,
  phi = 1,
  cutpoints = c(-1, 1),
  beta_coef = NULL,
  beta_type = "continuous",
  treat_assign = 0.5,
  return_data = FALSE,
  seed = as.numeric(Sys.time()),
  ...
)

```

Arguments

N	The sample size for the simulation. Include a vector of integers to examine power/results for multiple sample sizes.
k	The number of covariates/predictors.
iter	The number of simulations to run. For power calculation, should be at least 500 (yes, this will take some time).
cores	The number of cores to use to parallelize the simulation.
phi	Value of the dispersion parameter in the beta distribution.
cutpoints	Value of the two cutpoints for the ordered model. By default are the values -1 and +1 (these are interpreted in the logit scale and so should not be too large). The farther apart, the fewer degenerate (0 or 1) responses there will be in the distribution.
beta_coef	If not null, a vector of length k of the true predictor coefficients/treatment values to use for the simulation. Otherwise, coefficients are drawn from a random uniform distribution from -1 to 1 for each predictor.
beta_type	Can be either continuous or binary. Use the latter for conventional treatments with two values.
treat_assign	If beta_type is set to binary, you can use this parameter to set the proportion of N assigned to treatment. By default, the parameter is set to 0.5 for equal/balanced treatment control groups.
return_data	Whether to return the simulated data as a list in the data column of the returned data frame.
seed	The seed to use to make the results reproducible. Set automatically to a date-time stamp.
...	Any other arguments are passed on to the brms::brm function to control modeling options.

Details

This function implements the simulation found in Kubinec (2022). This simulation allows you to vary the sample size, number & type of predictors, values of the predictors (or treatment values), and the power to target. The function returns a data frame with one row per simulation draw and covariate k .

Value

a tibble data frame with columns of simulated and estimated values and rows for each simulation iteration X coefficient combination. I.e., if there are five predictors, and 1,000 iterations, the resulting data frame will have 1,000 rows. If there are multiple values for N , then each value of N will have its own set of iterations, making the final size of the data a multiple of the number of sample sizes to iterate over. The data frame will have the following columns: 1.

Examples

```
# This function takes a while to run as it has
# to fit an ordered beta regression to each
# draw. The package comes with a saved
# simulation dataset you can inspect to see what the
# result looks like

data("sim_data")

library(dplyr)

# will take a while to run this

sim_data <- sim_ordbeta(N=c(250,750),
  k=1,
  beta_coef = .5,
  iter=5,cores=2,
  beta_type="binary",
  treat_assign=0.3)

# to get the power values by N, simply summarize/group
# by N with functions from the R package dplyr

sim_data %>%
  group_by(N) %>%
  summarize(mean_power=mean(power))
```

Index

* datasets

- fit_imputed, 3
- fit_multivariate, 3
- ord_fit_mean, 8
- ord_fit_phi, 9
- pew, 9
- sim_data, 12

- brms::bf, 5
- brms::brm, 7, 13
- brms::brm_multiple, 6
- brms::pp_check, 10
- brms::prior_string, 7
- brms::set_prior, 7

- dordbeta, 2

- fit_imputed, 3
- fit_multivariate, 3

- normalize, 4, 7

- ord_fit_mean, 8
- ord_fit_phi, 9
- ordbetareg, 5, 10

- pew, 9
- pp_check_ordbeta, 10

- rordbeta, 11

- sim_data, 12
- sim_ordbeta, 12