# The Beta-distribution

## Parametrisation

The Beta-distribution has the following density

$$\pi(y) = \frac{1}{B(a,b)} y^{a-1}(1-y)^{b-1}, \qquad 0 < y < 1, \quad a > 0, \quad b > 0$$

where $B(a,b)$ is the Beta-function

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

and $\Gamma(x)$ is the Gamma-function. The (re-)parameterisation used is

$$\mu = \frac{a}{a+b}, \qquad 0 < \mu < 1$$

and

$$\phi = a + b, \qquad \phi > 0,$$

as it makes

$$\mathrm{E}(y) = \mu \quad \text{and} \quad \mathrm{Var}(y) = \frac{\mu(1-\mu)}{1+\phi}.$$

The parameter $\phi$ is know as the *precision parameter*, since for fixed $\mu$, the larger $\phi$ the smaller the variance of $y$. The parameters $\{a,b\}$ are given as $\{\mu,\phi\}$ as follows,

$$a = \mu\phi \quad \text{and} \quad b = -\mu\phi + \phi.$$

In some applications then observations close to 0 or 1, are censored and represented as exactly 0 and 1. For this, we introduced a censor value $0 < \delta < 1/2$ and treat all $y \le \delta$ or $y \ge 1 - \delta$ as censored observations. By default, no censoring is applied ($\delta = 0$).

## Link-function

The linear predictor $\eta$ is linked to the mean $\mu$ using a default logit-link

$$\mu = \frac{\exp(\eta)}{1 + \exp(\eta)}.$$

## Hyperparameter

The hyperparameter is the precision parameter $\phi$, which is represented as

$$\phi = s_i \exp(\theta)$$

where $s = (s_i) > 0$ is a fixed scaling, and the prior is defined on $\theta$.

## Specification

- family = `beta`

- Required argument: $y$

- Optional argument: $s$ (argument `scale`, default all 1, $s > 0$)

- Optional argument: truncation limit $0 \le \delta < 1/2$ (argument `beta.trunctation`, $\delta = 0$ means no trunctation).

**Hyperparameter spesification and default values**

**doc** The Beta likelihood

**hyper**

> **theta**
>
>> **hyperid** 61001
>> **name** precision parameter
>> **short.name** phi
>> **initial** 2.30258509299405
>> **fixed** FALSE
>> **prior** loggamma
>> **param** 1 0.1
>> **to.theta** function(x) log(x)
>> **from.theta** function(x) exp(x)

**survival** FALSE

**discrete** FALSE

**link** default logit loga cauchit probit cloglog loglog

**pdf** beta

## Example

In the following example we estimate the parameters in a simulated example.

```
n = 1000
w = runif(n, min = 0.25,  max = 0.75)
phi = 5 * w
z = rnorm(n, sd=0.2)
eta = 1 + z
mu = exp(eta)/(1+exp(eta))
a = mu * phi
b = -mu * phi + phi
y = rbeta(n, a, b)


formula = y ~ 1 + z
r = inla(formula, data = data.frame(y, z, w),
         family = "beta", scale = w)
summary(r)
```

In this example we add truncation.

```
## the precision parameter in the beta distribution
phi = 5

## generate simulated data
n = 1000
z = rnorm(n, sd=.2)
eta = 1 + z
```

```
mu = exp(eta)/(1+exp(eta))
a = mu * phi
b = -mu * phi + phi
y = rbeta(n, a, b)

## this is the censoring
cens <- 0.05
y[y <= cens] <- 0
y[y >= 1-cens] <- 1

## estimate the model
formula = y ~ 1 + z
r = inla(formula, data = data.frame(y, z), family = "beta",
         control.family = list(beta.censor.value = cens))
summary(r)
```

## Notes

None.