

Skew-Normal stochastic volatility likelihood

Parametrisation

The standardised Skew-Normal distribution is

$$f(z) = \frac{2}{\omega_\alpha} \phi\left(\frac{z - \xi_\alpha}{\omega_\alpha}\right) \Phi\left(\alpha \frac{z - \xi_\alpha}{\omega_\alpha}\right)$$

where ω_α and ξ_α are so that the mean is zero and variance is one, and they are functions of the skewness parameter α . The skewness s is the standardised skewness (standardized third central moment), which is a function of α

The skew-normal stochastic likelihood is defined as the density wrt y , where y is Skew-Normal distributed with zero mean, skewness s and variance

$$\text{variance} = \exp(\eta) + 1/\tau$$

and

η : is the linear predictor

τ : is an offset in the variance

Link-function

The variance depends on the linear predictor

$$\mu = \exp(\eta) + 1/\tau$$

Hyperparameters

The (standardised) skewness s , is represented as

$$\gamma = 0.988(2 \frac{\exp(\theta_1)}{1 + \exp(\theta_1)} - 1)$$

and the prior is defined on θ_1 .

The offset in the variance, $1/\tau$ is represented as

$$\tau = \exp(\theta_2)$$

and the prior is defined on θ_2 . (By default θ_2 is fixed to a high value which makes $1/\tau \approx 0$.)

Specification

- family = `stochvol`
- Required arguments: y .

Hyperparameter specification and default values

doc The SkewNormal stochvol likelihood

hyper

theta1

hyperid 82101

```

    name logit skew
    short.name skew
    initial 0.00123456789
    fixed FALSE
    prior pc.sn
    param 10
    to.theta function(x, skew.max = 0.988) log((1 + x / skew.max) / (1 - x / skew.max))
    from.theta function(x, skew.max = 0.988) skew.max * (2 * exp(x) / (1 + exp(x)) - 1)
  theta2
    hyperid 82102
    name log precision
    short.name prec
    initial 500
    fixed TRUE
    prior loggamma
    param 1 0.005
    to.theta function(x) log(x)
    from.theta function(x) exp(x)

status experimental

survival FALSE

discrete FALSE

link default log

pdf stochvolns

```

Example

```

library(sn)
n <- 1000
x <- scale(arima.sim(n, model= list(ar = 0.95)))
skew <- 0.2
y <- numeric(n)
for(i in 1:n) {
  variance <- exp(x[i])
  par <- unlist(INLA::inla.sn.reparam(moments = c(0, variance, skew)))
  y[i] <- rsn(1, dp = par)
}

r = inla(y ~ 1 + f(idx, model="ar1",
  hyper = list(
    prec = list(prior = "pc.prec",
      param = c(0.5, 0.01)),
    rho = list(prior = "pc.cor1",
      param = c(0.8, 0.5))),
  control.fixed = list(prec.intercept = 1),
  data = data.frame(y, idx=1:n),

```

```
family = "stochvolnsn",
control.inla = list(cmin = 0, b.strategy="skip"),
num.threads = "3:1",
verbose = TRUE)
summary(r)
```

Notes

This implementation is similar to `family="sn"`, see also that documentation.