

The gammacountmean likelihood

Parametrisation

The `gammacount_mu` distribution is a discrete probability distribution on $0, 1, 2, 3, \dots$, derived from a renewal process with Gamma-distributed inter-arrival times. The probability mass function is

$$\text{Prob}(y) = G(T; y\alpha, \beta) - G(T; (y+1)\alpha, \beta)$$

where $G(\cdot; a, b)$ is the cumulative distribution function of the Gamma distribution, with shape $a > 0$ and rate $b > 0$,

$$G(T; a, b) = \frac{b^a}{\Gamma(a)} \int_0^T u^{a-1} e^{-bu} du.$$

This model exhibits:

- underdispersion if $\alpha > 1$ ($\text{Var}(Y) < \mu$),
- equidispersion if $\alpha = 1$ (Poisson limit),
- overdispersion if $0 < \alpha < 1$ ($\text{Var}(Y) > \mu$).

Link-function

The linear predictor η is linked to the event rate $\lambda = \mathbb{E}[Y]/T$ via

$$\log(\lambda) = \eta,$$

and the Gamma rate parameter β is related to λ through the mean-parametrized renewal link

$$\beta = \alpha E \lambda + \frac{1}{2T}.$$

This formulation accommodates unequal observation windows T_i and provides direct interpretation of $E \exp(\eta_j)$ as the multiplicative change in event rate per unit time. Here, E typically is (population or at-risk size, and default equal to 1 if not given).

Hyperparameter

The dispersion parameter $\alpha > 0$ is represented

$$\alpha = \exp(\theta),$$

with a prior defined on the log-scale parameter θ .

Specification

- `family = "gammacountmean"`
- Required arguments: response y and exposure time T , provided as `inla.mdata(y, T)` or `inla.mdata(y, T, E)` (see example).
- If E is not given, then $E = 1$.

Hyperparameter specification and default values

doc Another Gamma generalisation of the Poisson likelihood

hyper

theta

```
hyperid 59101
name log alpha
short.name alpha
output.name Log-alpha parameter for Gammacountmean observations
output.name.intern Alpha parameter for Gammacountmean observations
initial 0
fixed FALSE
prior pc.gammacount
param 3
to.theta function(x) log(x)
from.theta function(x) exp(x)
```

survival FALSE

discrete FALSE

link default log

pdf gammacountmean

Example

In the following example we estimate the parameters in a simulated dataset with varying exposure times.

```
## =====
## Rate-Based Mean-Parametrized Gamma-Count Regression
## Models:  $\log(\lambda_i) = \log(E) + x_i' \gamma$ , where
##  $\lambda_i = E[N(T_i)] / T_i$ 
## =====

## -----
## Exact renewal-process simulator
## -----

rGC_renewal <- function(n, alpha, beta, T) {
  y <- integer(n)
  for (i in seq_len(n)) {
    t <- 0
    k <- 0L
    repeat {
      t <- t + rgamma(1, shape = alpha, rate = beta[i])
      if (t > T[i]) break
      k <- k + 1L
    }
    y[i] <- k
  }
}
```

```

    }
  y
}

## -----
## Exact Gamma-count PMF (vectorized over T)
## -----
dGC_renewal <- function(y, alpha, beta, T, log = FALSE) {
  Fy1 <- pgamma(T, shape = (y + 1) * alpha, rate = beta)
  Fy0 <- ifelse(y == 0, 1, pgamma(T, shape = y * alpha, rate = beta))
  p <- Fy0 - Fy1
  p <- pmax(p, 1e-14)
  if (log) log(p) else p
}

## -----
## Negative log-likelihood
## -----
negloglik_gc_rate <- function(par, y, X, T, E) {
  if (missing(E))
    E <- rep(1, length(T))

  log_alpha <- par[1]
  gamma <- par[-1]

  alpha <- exp(log_alpha)
  if (!is.finite(alpha) || alpha < 0.05) return(1e12)

  ## Model log(rate): lambda = exp(X %*% gamma)
  eta <- X %*% gamma
  lambda <- as.numeric(E * exp(eta))
  if (any(!is.finite(lambda)) || any(lambda < 1e-6)) return(1e12)

  ## Mean count: mu = lambda * T
  mu <- lambda * T

  ## Link: beta = (alpha * mu + 0.5) / T = alpha * lambda + 0.5 / T
  beta <- alpha * lambda + 0.5 / T
  if (any(beta <= 0)) return(1e12)

  ll <- sum(dGC_renewal(y, alpha, beta, T, log = TRUE))
  if (!is.finite(ll)) return(1e12)

  -ll
}

## -----
## Simulate data
## -----
n <- 300
x <- rnorm(n, sd = 0.3)

```

```

X <- cbind(1, x)
E <- runif(n)

## Varying observation windows
T_obs <- runif(n, 0.5, 3) # e.g., different durations per unit

alpha_true <- 2.0
gamma_true <- c(1.0, -0.6) # log(rate) = gamma0 + gamma1 * x

lambda_true <- as.numeric(E * exp(X %*% gamma_true)) # true rate
mu_true <- lambda_true * T_obs # expected count
beta_true <- alpha_true * lambda_true + 0.5 / T_obs

y <- rGC_renewal(n, alpha_true, beta_true, T_obs)

## -----
## ML estimation
## -----
## Initialize with Poisson log-rate
glm_init <- glm(y ~ x, family = poisson(), offset = log(E) + log(T_obs))
gamma_init <- coef(glm_init)
par_init <- c(log(1), gamma_init)

fit <- optim(
  par = par_init,
  fn = negloglik_gc_rate,
  y = y, X = X, T = T_obs, E = E,
  method = "BFGS", control = list(maxit = 1000, reltol = 1e-10))
alpha_est <- exp(fit$par[1])
gamma_est <- fit$par[-1]

r <- inla(Y ~ 1 + x,
  family = "gammacountmean",
  data = list(Y = inla.mdata(y, T_obs, E), x = x))
print(c(alpha_est, gamma_est))
summary(r)

```

Notes

- This model extends the classical Gamma-count framework (Winkelmann, 1995; Zeviani et al., 2014) by directly regressing on the event rate λ , unlike the original `gammacount` family in R-INLA (Nadifar et al., 2023), which assumes fixed $T = 1$ and models waiting times.
- The approximation $E[Y] \approx \beta T / \alpha - 1 / (2\alpha)$ is asymptotically valid for moderate-to-large T (Smith, 1958; Cox, 1962).

References

- Cox, D. R. (1962). *Renewal Theory*. Methuen.
- Nadifar, M., Baghishani, H., Fallah, A. (2023). A flexible generalized Poisson likelihood for spatial counts constructed by renewal theory, motivated by groundwater quality assessment.

Journal of Agricultural, Biological, and Environmental Statistics, 28(4), 726-748.

- Smith, W. L. (1958). Renewal theory and its ramifications. *Journal of the Royal Statistical Society: Series B*, 20(2), 243–302.
- Winkelmann, R. (1995). Duration dependence and dispersion in count data models. *Journal of Business & Economic Statistics*, 13(4), 467–474.
- Zeviani, W. M., et al. (2014). The Gamma-count distribution in the analysis of experimental underdispersed data. *Journal of Applied Statistics*, 41(7), 1513–1527.