Description

bayes: xtnbreg fits a Bayesian panel-data random-effects negative binomial model to a nonnegative count outcome; see [BAYES] **bayes** and [XT] **xtnbreg** for details.

Quick start

Bayesian random-effects negative binomial model of y on x1 and x2 with random intercepts by id (after xtsetting on panel variable id), using default normal priors for regression coefficients and beta prior for the random effects, and Pareto prior for the shape parameters of the beta prior

    bayes: xtnbreg y x1 x2

Use a standard deviation of 10 instead of 100 for the default normal priors

    bayes, normalprior(10): xtnbreg y x1 x2

Use uniform priors for the slopes and a normal prior for the intercept

    bayes, prior({y: x1 x2}, uniform(-10,10)) ///
        prior({y:_cons}, normal(0,10)): xtnbreg y x1 x2

Save simulation results to simdata.dta, and use a random-number seed for reproducibility

    bayes, saving(simdata) rseed(123): xtnbreg y x1 x2

Specify 20,000 Markov chain Monte Carlo (MCMC) samples, set length of the burn-in period to 5,000, and request that a dot be displayed every 500 simulations

    bayes, mcmcsize(20000) burnin(5000) dots(500): xtnbreg y x1 x2

In the above, request that the 90% highest posterior density (HPD) credible interval be displayed instead of the default 95% equal-tailed credible interval

    bayes, clevel(90) hpd

Display incidence-rate ratios instead of coefficients

    bayes: xtnbreg y x1 x2, irr

Display incidence-rate ratios on replay

    bayes, irr

Random-effects model with exposure variable evar

    bayes: xtnbreg y x1 x2, exposure(evar)

Also see *Quick start* in [BAYES] **bayes** and *Quick start* in [XT] **xtnbreg**.

Menu

Statistics > Longitudinal/panel data > Bayesian regression > Negative binomial regression

Syntax

bayes [, *bayesopts*] : xtnbreg *depvar* [*indepvars*] [*if*] [*in*] [, *options*]

options	Description
Model	
noconstant	suppress constant term
exposure(*varname_e*)	include ln(*varname_e*) in model with coefficient constrained to 1
offset(*varname_o*)	include *varname_o* in model with coefficient constrained to 1
Reporting	
irr	report incidence-rate ratios
display_options	control spacing, line width, and base and empty cells
level(*#*)	set credible level; default is level(95)

A panel variable must be specified; see [XT] **xtset**.

indepvars may contain factor variables; see [U] **11.4.3 Factor variables**.

depvar, *indepvars*, *varname_e*, and *varname_o* may contain time-series operators; see [U] **11.4.4 Time-series varlists**.

bayes: xtnbreg, level() is equivalent to bayes, clevel(): xtnbreg.

For a detailed description of options, see *Options for RE/FE models* in [XT] **xtnbreg**.

bayesopts	Description
Priors	
* normalprior(*#*)	specify standard deviation of default normal priors for regression coefficients; default is normalprior(100)
prior(*priorspec*)	prior for model parameters; this option may be repeated
dryrun	show model summary without estimation
Simulation	
nchains(*#*)	number of chains; default is to simulate one chain
mcmcsize(*#*)	MCMC sample size; default is mcmcsize(10000)
burnin(*#*)	burn-in period; default is burnin(2500)
thinning(*#*)	thinning interval; default is thinning(1)
rseed(*#*)	random-number seed
exclude(*paramref*)	specify model parameters to be excluded from the simulation results
Blocking	
block(*paramref* [, *blockopts*])	specify a block of model parameters; this option may be repeated
blocksummary	display block summary
Initialization	
initial(*initspec*)	specify initial values for model parameters with a single chain
init*#*(*initspec*)	specify initial values for #th chain; requires nchains()
initall(*initspec*)	specify initial values for all chains; requires nchains()
nomleinitial	suppress the use of maximum likelihood estimates as starting values
initrandom	specify random initial values
initsummary	display initial values used for simulation
* noisily	display output from the estimation command during initialization

[Adaptation]
  <u>adapt</u>ation(*adaptopts*)   control the adaptive MCMC procedure
  <u>sc</u>ale(*#*)   initial multiplier for scale factor; default is scale(2.38)
  <u>cov</u>ariance(*cov*)   initial proposal covariance; default is the identity matrix

[Reporting]
  <u>clev</u>el(*#*)   set credible interval level; default is clevel(95)
  hpd   display HPD credible intervals instead of the default equal-tailed
                                 credible intervals
 * <u>irr</u>   report incidence-rate ratios
  <u>ef</u>orm[(*string*)]   report exponentiated coefficients and, optionally, label as *string*
  remargl   compute log marginal-likelihood; suppressed by default
  batch(*#*)   specify length of block for batch-means calculations;
                                 default is batch(0)
  <u>saving</u>(*filename*[, replace]) save simulation results to *filename*.dta
  <u>nomodelsummary</u>   suppress model summary
  chainsdetail   display detailed simulation summary for each chain
  [no]dots   suppress dots or display dots every 100 iterations and iteration
                                 numbers every 1,000 iterations; default is nodots
  dots(*#*[, every(*#*)])   display dots as simulation is performed
  [no]show(*paramref*)   specify model parameters to be excluded from or included in
                                 the output
  <u>showreffects</u>[(*reref*)]   specify that all or a subset of random-effects parameters be included
                                 in the output
  <u>notable</u>   suppress estimation table
  <u>noheader</u>   suppress output header
  title(*string*)   display *string* as title above the table of parameter estimates
  display_options   control spacing, line width, and base and empty cells

[Advanced]
  search(*search_options*)   control the search for feasible initial values
  corrlag(*#*)   specify maximum autocorrelation lag; default varies
  corrtol(*#*)   specify autocorrelation tolerance; default is corrtol(0.01)

---

* Starred options are specific to the bayes prefix; other options are common between bayes and bayesmh.

Options prior() and block() may be repeated.

priorspec and *paramref* are defined in [BAYES] **bayesmh**.

paramref may contain factor variables; see [U] **11.4.3 Factor variables**.

collect is allowed; see [U] **11.1.10 Prefix commands**.

See [U] **20 Estimation and postestimation commands** for more capabilities of estimation commands.

Model parameters are regression coefficients {*depvar*:*indepvars*}, random effects {U[*panelvar*]} or simply {U}, and shape
   parameters {r} and {s} for the beta prior of {U}; also see *Methods and formulas*. Use the dryrun option to see the
   definitions of model parameters prior to estimation.

For a detailed description of *bayesopts*, see *Options* in [BAYES] **bayes**.

Remarks and examples

For a general introduction to Bayesian analysis, see [BAYES] **Intro**. For a general introduction to
Bayesian estimation using an adaptive Metropolis–Hastings algorithm, see [BAYES] **bayesmh**. For re-
marks and examples specific to the bayes prefix, see [BAYES] **bayes**. For details about the estimation
command, see [XT] **xtnbreg**.

For a simple example of the bayes prefix, see *Introductory example* in [BAYES] **bayes**. Also see *Panel-data models* in [BAYES] **bayes**.

Stored results

See *Stored results* in [BAYES] **bayes**. In addition, bayes: xtnbreg also stores the following results:

Macros
    e(ivar)   variable denoting groups
    e(redistrib)   distribution of random effects

Methods and formulas

Bayesian random-effects negative binomial models are based on random-effects negative binomial models described in *Methods and formulas* of [XT] **xtnbreg**.

Let y_{it} be the count for the tth observation in the ith group. We assume $y_{it} \mid \gamma_{it} \sim \text{Poisson}(\gamma_{it})$, where $\gamma_{it} \mid u_i \sim \text{gamma}\{\lambda_{it}, (1-u_i)/u_i\}$ with $\lambda_{it} = \exp(\mathbf{x}_{it}\boldsymbol{\beta} + \text{offset}_{it})$ and u_i is a dispersion-control parameter that varies randomly across groups. The likelihood of the model is thus

$$\Pr(Y_{it} = y_{it} \mid \mathbf{x}_{it}, u_i) = \frac{\Gamma(\lambda_{it} + y_{it})}{\Gamma(\lambda_{it})\Gamma(y_{it} + 1)} u_i^{\lambda_{it}} (1 - u_i)^{y_{it}}$$

We further assume that random-effects dispersion parameters u_i's are a priori independent and follow beta distribution with shape parameters r and s, $u_i \sim \text{Beta}(r, s)$. The hyperprior for the shape parameters is chosen so that the joint distribution of the mean $r/(r+s)$ and inverse square-root of the sample size $(r+s)$ of the beta distribution is uniform, $\{r/(r+s), (r+s)^{-0.5}\} \sim 1$; see Gelman et al. (2014, sec. 5.3). This choice leads to a diffused distribution for (r, s) with a density proportional to $(r+s)^{-2.5}$, which is a Pareto-type distribution.

bayes: xtnbreg uses the default initial value of 0.5 for the random effects u_i's.

See *Methods and formulas* in [BAYES] **bayesmh**.

Reference

Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin. 2014. *Bayesian Data Analysis*. 3rd ed. Boca Raton, FL: Chapman and Hall/CRC.

Also see

[BAYES] **bayes** — Bayesian regression models using the bayes prefix

[XT] **xtnbreg** — Fixed-effects, random-effects, & population-averaged negative binomial models

[BAYES] **Bayesian postestimation** — Postestimation tools after Bayesian estimation

[BAYES] **Bayesian estimation** — Bayesian estimation commands

[BAYES] **Bayesian commands** — Introduction to commands for Bayesian analysis

[BAYES] **Intro** — Introduction to Bayesian analysis

[BAYES] **Glossary**

For suggested citations, see the FAQ on citing Stata documentation.