

[Description](#)[Remarks and examples](#)[Quick start](#)[Stored results](#)[Menu](#)[Methods and formulas](#)[Syntax](#)[Also see](#)

Description

`bayes: zinb` fits a Bayesian zero-inflated negative binomial regression to a nonnegative count outcome with a high fraction of zeros; see [\[BAYES\] bayes](#) and [\[R\] zinb](#) for details.

Quick start

Bayesian zero-inflated negative binomial regression of y on x_1 and x_2 , using z to model excess zeros and using default normal priors for regression coefficients and log-overdispersion parameter

```
bayes: zinb y x1 x2, inflate(z)
```

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

```
bayes, normalprior(10): zinb y x1 x2, inflate(z)
```

Use uniform priors for the slopes and a normal prior for the intercept of the main regression

```
bayes, prior({y: x1 x2}, uniform(-10,10)) ///  
prior({y: _cons}, normal(0,10)): zinb y x1 x2, inflate(z)
```

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

```
bayes, saving(simdata) rseed(123): zinb y x1 x2, inflate(z)
```

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, mcmcsample(20000) burnin(5000) dots(500): zinb y x1 x2, inflate(z)
```

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: zinb y x1 x2, inflate(z) irr
```

Display incidence-rate ratios on replay

```
bayes, irr
```

Also see [Quick start](#) in [\[BAYES\] bayes](#) and [Quick start](#) in [\[R\] zinb](#).

Menu

Statistics > Count outcomes > Bayesian regression > Zero-inflated negative binomial regression

Syntax

```
bayes [ , bayesopts ] : zinb depvar [indepvars] [if] [in] [weight] ,
  inflate(varlist [ , offset(varname) ] | _cons) [options]
```

<i>options</i>	Description
----------------	-------------

Model	
* <u>inflate</u> ()	equation that determines whether the count is zero
<u>noconstant</u>	suppress constant term
<u>exposure</u> (<i>varname</i> _e)	include ln(<i>varname</i> _e) in model with coefficient constrained to 1
<u>offset</u> (<i>varname</i> _o)	include <i>varname</i> _o in model with coefficient constrained to 1
<u>probit</u>	use probit model to characterize excess zeros; default is logit

Reporting	
<u>irr</u>	report incidence-rate ratios
<u>display_options</u>	control spacing, line width, and base and empty cells
<u>level</u> (#)	set credible level; default is level(95)

*inflate(*varlist*[, offset(*varname*)] | _cons) is required.

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

fweights are allowed; see [U] 11.1.6 weight.

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

For a detailed description of *options*, see *Options* in [R] zinb.

<i>bayesopts</i>	Description
------------------	-------------

Priors	
* <u>normalprior</u> (#)	specify standard deviation of default normal priors for regression coefficients and log-overdispersion parameter; default is normalprior(100)
<u>prior</u> (<i>priorspec</i>)	prior for model parameters; this option may be repeated
<u>dryrun</u>	show model summary without estimation

Simulation	
<u>nchains</u> (#)	number of chains; default is to simulate one chain
<u>mcmcsize</u> (#)	MCMC sample size; default is mcmcsize(10000)
<u>burnin</u> (#)	burn-in period; default is burnin(2500)
<u>thinning</u> (#)	thinning interval; default is thinning(1)
<u>rseed</u> (#)	random-number seed
<u>exclude</u> (<i>paramref</i>)	specify model parameters to be excluded from the simulation results

Blocking	
* <u>blocksize</u> (#)	maximum block size; default is blocksize(50)
<u>block</u> (<i>paramref</i> [, <i>blockopts</i>])	specify a block of model parameters; this option may be repeated
<u>blocksummary</u>	display block summary
* <u>noblocking</u>	do not block parameters by default

Initialization

<code>initial(<i>initspec</i>)</code>	specify initial values for model parameters with a single chain
<code>init#(<i>initspec</i>)</code>	specify initial values for #th chain; requires <code>nchains()</code>
<code>initall(<i>initspec</i>)</code>	specify initial values for all chains; requires <code>nchains()</code>
<code>nomleinitial</code>	suppress the use of maximum likelihood estimates as starting values
<code>initrandom</code>	specify random initial values
<code>initsummary</code>	display initial values used for simulation
* <code>noisily</code>	display output from the estimation command during initialization

Adaptation

<code>adaptation(<i>adaptopts</i>)</code>	control the adaptive MCMC procedure
<code>scale(#)</code>	initial multiplier for scale factor; default is <code>scale(2.38)</code>
<code>covariance(<i>cov</i>)</code>	initial proposal covariance; default is the identity matrix

Reporting

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

Advanced

<code>search(<i>search_options</i>)</code>	control the search for feasible initial values
<code>corrlag(#)</code>	specify maximum autocorrelation lag; default varies
<code>corrtol(#)</code>	specify autocorrelation tolerance; default is <code>corrtol(0.01)</code>

* 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}` for the main regression and `{inflate: varlist}` for the inflation equation and log-overdispersion parameter `{lnalpha}`. 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 remarks and examples specific to the bayes prefix, see [BAYES] **bayes**. For details about the estimation command, see [R] **zinb**.

For a simple example of the bayes prefix, see *Introductory example* in [BAYES] **bayes**. Also see *Zero-inflated negative binomial model* in [BAYES] **bayes**.

Stored results

See *Stored results* in [BAYES] **bayes**.

Methods and formulas

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

Also see

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

[R] **zinb** — Zero-inflated negative binomial regression

[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**

Stata, Stata Press, and Mata are registered trademarks of StataCorp LLC. Stata and Stata Press are registered trademarks with the World Intellectual Property Organization of the United Nations. StataNow and NetCourseNow are trademarks of StataCorp LLC. Other brand and product names are registered trademarks or trademarks of their respective companies. Copyright © 1985–2025 StataCorp LLC, College Station, TX, USA. All rights reserved.



For suggested citations, see the FAQ on [citing Stata documentation](#).