

bayes: mixed — Bayesian multilevel linear regression
[Description](#)[Remarks and examples](#)[Quick start](#)[Stored results](#)[Menu](#)[Methods and formulas](#)[Syntax](#)[Also see](#)

Description

`bayes: mixed` fits a Bayesian multilevel linear regression to a continuous outcome; see [\[BAYES\] bayes](#) and [\[ME\] mixed](#) for details.

Quick start

Bayesian two-level linear regression of `y` on `x1` and `x2` with random intercepts by `id`, using default normal priors for regression coefficients and default inverse-gamma priors for the error variance and for the variance of random intercepts

```
bayes: mixed y x1 x2 || id:
```

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

```
bayes, normalprior(10): mixed y x1 x2 || id:
```

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)): mixed y x1 x2 || id:
```

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

```
bayes, saving(simdata) rseed(123): mixed y x1 x2 || id:
```

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): mixed y x1 x2 || id:
```

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
```

Also see [Quick start](#) in [\[BAYES\] bayes](#) and [Quick start](#) in [\[ME\] mixed](#).

Menu

Statistics > Multilevel mixed-effects models > Bayesian regression > Linear regression

Syntax

```
bayes [ , bayesopts ] : mixed depvar fe_equation
      [ || re_equation ] [ || re_equation ... ] [ , options ]
```

where the syntax of *fe_equation* is

```
[ indepvars ] [ if ] [ in ] [ weight ] [ , fe_options ]
```

and the syntax of *re_equation* is one of the following:

for random coefficients and intercepts

```
levelvar: [ varlist ] [ , re_options ]
```

for random effects among the values of a factor variable

```
levelvar: R.varname
```

levelvar either is a variable identifying the group structure for the random effects at that level or is `_all`, representing one group comprising all observations.

<i>fe_options</i>	Description
-------------------	-------------

Model

<code><u>noconstant</u></code>	suppress constant term from the fixed-effects equation
--------------------------------	--

<i>re_options</i>	Description
-------------------	-------------

Model

<code><u>covariance</u>(<i>vartype</i>)</code>	variance–covariance structure of the random effects ; structures independent , exchangeable , identity , and unstructured are supported
<code><u>noconstant</u></code>	suppress constant term from the random-effects equation

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

Reporting

<code><u>noheader</u></code>	suppress output header
<code><u>nogroup</u></code>	suppress table summarizing groups
<code><i>display_options</i></code>	control spacing, line width, and base and empty cells
<code><u>level</u>(#)</code>	set credible level; default is <code>level(95)</code>

indepvars may contain factor variables; see [\[U\] 11.4.3 Factor variables](#).

depvar, *indepvars*, and *varlist* may contain time-series operators; see [\[U\] 11.4.4 Time-series varlists](#).

fweights are allowed; see [\[U\] 11.1.6 weight](#).

`bayes: mixed`, `level()` is equivalent to `bayes, clevel(): mixed`.

For a detailed description of *options*, see [Options](#) in [\[ME\] mixed](#).

<i>bayesopts</i>	Description
Priors	
* <u>normalprior</u> (#)	specify standard deviation of default normal priors for regression coefficients; default is <code>normalprior(100)</code>
* <u>igammaprior</u> (# #)	specify shape and scale of default inverse-gamma prior for variance components; default is <code>igammaprior(0.01 0.01)</code>
* <u>wishartprior</u> (# [...])	specify degrees of freedom and, optionally, scale matrix of default inverse-Wishart prior for unstructured random-effects covariance
<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 <code>mcmcsize(10000)</code>
<u>burnin</u> (#)	burn-in period; default is <code>burnin(2500)</code>
<u>thinning</u> (#)	thinning interval; default is <code>thinning(1)</code>
<u>rseed</u> (#)	random-number seed
<u>exclude</u> (<i>paramref</i>)	specify model parameters to be excluded from the simulation results
<u>restubs</u> (<i>restub1 restub2 ...</i>)	specify stubs for random-effects parameters for all levels
Blocking	
* <u>blocksize</u> (#)	maximum block size; default is <code>blocksize(50)</code>
<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	
<u>initial</u> (<i>initspec</i>)	specify initial values for model parameters with a single chain
<u>init#</u> (<i>initspec</i>)	specify initial values for #th chain; requires <code>nchains()</code>
<u>initall</u> (<i>initspec</i>)	specify initial values for all chains; requires <code>nchains()</code>
<u>nomleinitial</u>	suppress the use of maximum likelihood estimates as starting values
<u>initrandom</u>	specify random initial values
<u>initsummary</u>	display initial values used for simulation
* <u>noisily</u>	display output from the estimation command during initialization
Adaptation	
<u>adaptation</u> (<i>adaptopts</i>)	control the adaptive MCMC procedure
<u>scale</u> (#)	initial multiplier for scale factor; default is <code>scale(2.38)</code>
<u>covariance</u> (<i>cov</i>)	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>eform[<i>(string)</i>]</code>	report exponentiated coefficients and, optionally, label as <i>string</i>
<code>remargl</code>	compute log marginal-likelihood
<code>batch(#)</code>	specify length of block for batch-means calculations; default is <code>batch(0)</code>
<code>saving(filename[, replace])</code>	save simulation results to <i>filename.dta</i>
<code>nomodelsummary</code>	suppress model summary
<code>nomesummary</code>	suppress multilevel-structure 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>dots</code>
<code>dots(#[, every(#)])</code>	display dots as simulation is performed
<code>[no]show(paramref)</code>	specify model parameters to be excluded from or included in the output
<code>showeffects[<i>(ref)</i>]</code>	specify that all or a subset of random-effects parameters be included in the output
<code>melabel</code>	display estimation table using the same row labels as <code>mixed</code>
<code>nogroup</code>	suppress table summarizing groups
<code>notable</code>	suppress estimation table
<code>noheader</code>	suppress output header
<code>title(string)</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(search_options)</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}`, error variance `{e.depvar:sigma2}`, random effects `{rename}`, and either variance components `{rename:sigma2}` or, if option `covariance(unstructured)` is specified, matrix parameter `{restub:Sigma,matrix}`; see *Likelihood model* in [BAYES] `bayes` for how `renames` and `restub` are defined. 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

stata.com

For a general introduction to Bayesian analysis, see [BAYES] **Intro**. For a general introduction to Bayesian estimation using adaptive Metropolis–Hastings and Gibbs algorithms, see [BAYES] `bayesmh`. For remarks and examples specific to the `bayes` prefix, see [BAYES] `bayes`. For details about the estimation command, see [ME] `mixed`.

For a simple example of the `bayes` prefix, see *Introductory example* in [BAYES] `bayes`. For multilevel examples, see *Multilevel models* in [BAYES] `bayes`.

By default, `bayes: mixed` uses Gibbs sampling for all model parameters except the random-effects parameters. If you specify a `prior()` distribution for which Gibbs sampling is not available, `bayes: mixed` will switch to adaptive Metropolis–Hastings sampling. In general, `bayes: mixed` will try to use a more efficient Gibbs sampling for the model parameters whenever available.

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

[ME] `mixed` — Multilevel mixed-effects linear regression

[BAYES] `Bayesian postestimation` — Postestimation tools for `bayesmh` and the `bayes` prefix

[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. Other brand and product names are registered trademarks or trademarks of their respective companies. Copyright © 1985–2023 StataCorp LLC, College Station, TX, USA. All rights reserved.

