

bayesgraph — Graphical summaries and convergence diagnostics

[Description](#)
[Options](#)
[Also see](#)

[Quick start](#)
[Remarks and examples](#)

[Menu](#)
[Methods and formulas](#)

[Syntax](#)
[Reference](#)

Description

`bayesgraph` provides graphical summaries and convergence diagnostics for simulated posterior distributions (MCMC samples) of model parameters and functions of model parameters obtained from the `bayesmh` command. Graphical summaries include trace plots, autocorrelation plots, and various distributional plots.

Quick start

Trace plot, histogram, autocorrelation plot, and density plot for parameter `{p}`

```
bayesgraph diagnostics {p}
```

Add plots for parameter `{y:x1}`

```
bayesgraph diagnostics {p} {y:x1}
```

As above, but for all model parameters

```
bayesgraph diagnostics _all
```

As above, but for a function of model parameters `{y:x1}` and `{p}`

```
bayesgraph diagnostics ({y:x1}/{p})
```

Specify a blue trace plot line for all plots

```
bayesgraph diagnostics {p} {y:x1} {y:x2}, traceopts(lcolor(blue))
```

Specify a blue trace plot line only for the second trace plot

```
bayesgraph diagnostics {p} {y:x1} {y:x2}, trace2opts(lcolor(blue))
```

Trace plots for all parameters in a single graph

```
bayesgraph trace _all, byparm
```

Cumulative sum plot for parameter `{p}`

```
bayesgraph cusum {p}
```

Scatterplot matrix for parameters `{p}` and `{y:x1}`

```
bayesgraph matrix {p} {y:x1}
```

Autocorrelation plots for elements 1,1 and 2,1 of matrix parameter `{S}`

```
bayesgraph ac {S_1_1} {S_2_1}
```

Diagnostic plots for all parameters in the model and pause at least 3 seconds before displaying the next graph

```
bayesgraph diagnostics _all, sleep(3)
```

As above, but pause until the user presses any key

```
bayesgraph diagnostics _all, wait
```

As above, but close the current Graph window when the next graph is displayed

```
bayesgraph diagnostics _all, close
```

Menu

Statistics > Bayesian analysis > Graphical summaries

Syntax

Graphical summaries and convergence diagnostics for a single parameter

```
bayesgraph graph scalar_param [, singleopts]
```

Graphical summaries and convergence diagnostics for multiple parameters

```
bayesgraph graph spec [spec ...] [, multiopts]
```

```
bayesgraph matrix spec spec [spec ...] [, singleopts]
```

Graphical summaries and convergence diagnostics for all parameters

```
bayesgraph graph _all [, multiopts]
```

<i>graph</i>	Description
diagnostics	multiple diagnostics in compact form
trace	trace plots
ac	autocorrelation plots
histogram	histograms
kdensity	density plots
cusum	cumulative sum plots
matrix	scatterplot matrix

`bayesgraph matrix` requires at least two parameters.

scalar_param is a [scalar model parameter](#) specified as `{param}` or `{eqname:param}` or an expression *exprspec* of scalar model parameters. Matrix model parameters are not allowed, but you may refer to their individual elements.

exprspec is an optionally labeled expression of model parameters specified in parentheses:

```
( [exprlabel: ] expr )
```

exprlabel is a valid Stata name, and *expr* is a scalar expression that may not contain matrix model parameters. See [Specifying functions of model parameters](#) in [\[BAYES\] bayesmh postestimation](#) for examples.

spec is either *scalar_param* or *exprspec*.

<i>singleopts</i>	Description
-------------------	-------------

Options

<code>skip(#)</code>	skip every # observations from the MCMC sample; default is <code>skip(0)</code>
<code>name(name, ...)</code>	specify name of graph
<code>saving(filename, ...)</code>	save graph in file
<i>graphopts</i>	graph-specific options

<i>multiopts</i>	Description
------------------	-------------

Options

<code>byparm[(<i>grbyparmopts</i>)]</code>	specify the display of plots on one graph; default is separate graph for each plot; not allowed with graphs <code>diagnostics</code> and <code>matrix</code> or with option <code>combine()</code>
<code>combine[(<i>grcombineopts</i>)]</code>	specify the display of plots on one graph; recommended when the number of parameters is large; not allowed with graphs <code>diagnostics</code> and <code>matrix</code> or with option <code>byparm()</code>
<code>sleep(#)</code>	pause for # seconds between multiple graphs; default is <code>sleep(0)</code>
<code>wait</code>	pause until the <code>—more—</code> condition is cleared
<code>[no]close</code>	(do not) close Graph windows when the next graph is displayed with multiple graphs; default is <code>noclose</code>
<code>skip(#)</code>	skip every # observations from the MCMC sample; default is <code>skip(0)</code>
<code>name(namespec, ...)</code>	specify names of graphs
<code>saving(filespec, ...)</code>	save graphs in files
<code>graphopts(<i>graphopts</i>)</code>	control the look of all graphs; not allowed with <code>byparm()</code>
<code>graph[#]opts(<i>graphopts</i>)</code>	control the look of #th graph; not allowed with <code>byparm()</code>
<i>graphopts</i>	equivalent to <code>graphopts(<i>graphopts</i>)</code> ; only one may be specified

<i>graphopts</i>	Description
------------------	-------------

<i>diagnosticopts</i>	options for bayesgraph <code>diagnostics</code>
<i>tslineopts</i>	options for bayesgraph <code>trace</code> and bayesgraph <code>cusum</code>
<i>acopts</i>	options for bayesgraph <code>ac</code>
<i>histopts</i>	options for bayesgraph <code>histogram</code>
<i>kdensityopts</i>	options for bayesgraph <code>kdensity</code>
<i>gmatrixopts</i>	options for bayesgraph <code>matrix</code>

<i>diagnosticsopts</i>	Description
<code>traceopts(<i>tslineopts</i>)</code>	affect rendition of all trace plots
<code>trace[#]opts(<i>tslineopts</i>)</code>	affect rendition of #th trace plot
<code>acopts(<i>acopts</i>)</code>	affect rendition of all autocorrelation plots
<code>ac[#]opts(<i>acopts</i>)</code>	affect rendition of #th autocorrelation plot
<code>histopts(<i>histopts</i>)</code>	affect rendition of all histogram plots
<code>hist[#]opts(<i>histopts</i>)</code>	affect rendition of #th histogram plot
<code>kdensopts(<i>kdensityopts</i>)</code>	affect rendition of all density plots
<code>kdens[#]opts(<i>kdensityopts</i>)</code>	affect rendition of #th density plot
<i>grcombineopts</i>	any option documented in [G-2] graph combine

<i>acopts</i>	Description
<code>ci</code>	plot autocorrelations with confidence intervals; not allowed with <code>byparm()</code>
<i>acopts</i>	any options other than <code>generate()</code> documented for the <code>ac</code> command in [TS] corrgram

<i>kdensityopts</i>	Description
<i>kdensopts</i>	options for the overall kernel density plot
<code>show(<i>showspec</i>)</code>	show first-half density, second-half density, or both; default is <code>both</code>
<code>kdensfirst(<i>kdens1opts</i>)</code>	affect rendition of the first-half density plot
<code>kdenssecond(<i>kdens2opts</i>)</code>	affect rendition of the second-half density plot

Options

Options

`byparm[(grbyparmopts)]` specifies the display of all plots of parameters as subgraphs on one graph. By default, a separate graph is produced for each plot when multiple parameters are specified. This option is not allowed with `bayesgraph diagnostics` or `bayesgraph matrix` and may not be combined with option `combine()`. When many parameters or expressions are specified, this option may fail because of memory constraints. In that case, you may use option `combine()` instead.

grbyparmopts is any of the suboptions of `by()` documented in [G-3] *by_option*.

`byparm()` allows *y* scales to differ for all graph types and forces *x* scales to be the same only for `bayesgraph trace` and `bayesgraph cusum`. Use `noyrescale` within `byparm()` to specify a common *y* axis, and use `xrescale` or `noxrescale` to change the default behavior for the *x* axis.

`byparm()` with `bayesgraph trace` and `bayesgraph cusum` defaults to displaying multiple plots in one column to accommodate the *x* axis with many iterations. Use `norowcoldefault` within `byparm()` to switch back to the default behavior of options `rows()` and `cols()` of the [G-3] *by_option*.

`combine[(grcombineopts)]` specifies the display of all plots of parameters as subgraphs on one graph and is an alternative to `byparm()` with a large number of parameters. By default, a separate

graph is produced for each plot when multiple parameters are specified. This option is not allowed with `bayesgraph` `diagnostics` or `bayesgraph` `matrix` and may not be combined with option `byparm()`. It can be used in cases where a large number of parameters or expressions are specified and the `byparm()` option would cause an error because of memory constraints.

`grcombineopts` is any of the options documented in [G-2] [graph combine](#).

`sleep(#)` specifies pausing for `#` seconds before producing the next graph. This option is allowed only when multiple parameters are specified. This option may not be combined with `wait`, `combine()`, or `byparm()`.

`wait` causes `bayesgraph` to display `—more—` and pause until any key is pressed before producing the next graph. This option is allowed when multiple parameters are specified. This option may not be combined with `sleep()`, `combine()`, or `byparm()`. `wait` temporarily ignores the global setting that is specified using `set more off`.

`[no]close` specifies that, for multiple graphs, the Graph window be closed when the next graph is displayed. The default is `noclose` or to not close any Graph windows.

`skip(#)` specifies that every `#` observations from the MCMC sample not be used for computation. The default is `skip(0)` or to use all observations in the MCMC sample. Option `skip()` can be used to subsample or thin the chain. `skip(#)` is equivalent to a thinning interval of `#+1`. For example, if you specify `skip(1)`, corresponding to the thinning interval of 2, the command will skip every other observation in the sample and will use only observations 1, 3, 5, and so on in the computation. If you specify `skip(2)`, corresponding to the thinning interval of 3, the command will skip every 2 observations in the sample and will use only observations 1, 4, 7, and so on in the computation. `skip()` does not thin the chain in the sense of physically removing observations from the sample, as is done by `bayesmh`'s `thinning()` option. It only discards selected observations from the computation and leaves the original sample unmodified.

`name(namespec[, replace])` specifies the name of the graph or multiple graphs. See [G-3] [name_option](#) for a single graph. If multiple graphs are produced, then the argument of `name()` is either a list of names or a *stub*, in which case graphs are named `stub1`, `stub2`, and so on. With multiple graphs, if `name()` is not specified and neither `sleep()` nor `wait` is specified, `name(Graph_#)` is assumed, and thus the produced graphs may be replaced by subsequent `bayesgraph` commands.

The `replace` suboption causes existing graphs with the specified name or names to be replaced.

`saving(filespec[, replace])` specifies the filename or filenames to use to save the graph or multiple graphs to disk. See [G-3] [saving_option](#) for a single graph. If multiple graphs are produced, then the argument of `saving()` is either a list of filenames or a *stub*, in which case graphs are saved with filenames `stub1`, `stub2`, and so on.

The `replace` suboption specifies that the file (or files) may be replaced if it already exists.

`graphopts(graphopts)` and `graph[#]opts(graphopts)` affect the rendition of graphs. `graphopts()` affects the rendition of all graphs but may be overridden for specific graphs by using the `graph#opts()` option. The options specified within `graph#opts()` are specific for each type of graph.

The two specifications

```
bayesgraph ..., graphopts(graphopts)
```

and

```
bayesgraph ..., graphopts
```

are equivalent, but you may specify one or the other.

These options are not allowed with `byparm()` and when only one parameter is specified.

graphopts specifies options specific to each graph type.

diagnosticopts specifies options for use with `bayesgraph diagnostics`. See the corresponding table in the syntax diagram for a list of options.

tslineopts specifies options for use with `bayesgraph trace` and `bayesgraph cusum`. See the options of [TS] **tsline** except `by()`.

acopts specifies options for use with `bayesgraph ac`.

`ci` requests that the graph of autocorrelations with confidence intervals be plotted. By default, confidence intervals are not plotted. This option is not allowed with `byparm()`.

acopts specifies any options except `generate()` of the `ac` command in [TS] **corrgram**.

histopts specifies options for use with `bayesgraph histogram`. See options of [R] **histogram** except `by()`.

kdensityopts specifies options for use with `bayesgraph kdensity`.

kdensopts specifies options for the overall kernel density plot. See the options documented in [R] **kdensity** except `generate()` and `at()`.

`show(showspec)` specifies which kernel density curves to plot. *showspec* is one of `both`, `first`, `second`, or `none`. `show(both)`, the default, overlays both the first-half density curve and the second-half density curve with the overall kernel density curve. If `show(first)` is specified, only the first-half density curve, obtained from the first half of an MCMC sample, is plotted. If `show(second)` is specified, only the second-half density curve, obtained from the second half of an MCMC sample, is plotted. If `show(none)` is specified, only the overall kernel density curve is shown.

`kdensityfirst(kdens1opts)` specifies options of [G-2] **graph twoway kdensity** except `by()` to affect rendition of the first-half kernel density plot.

`kdensitysecond(kdens2opts)` specifies options of [G-2] **graph twoway kdensity** except `by()` to affect rendition of the second-half kernel density plot.

grmatrixopts specifies options for use with `bayesgraph matrix`. See the options of [G-2] **graph matrix** except `by()`.

Remarks and examples

[stata.com](http://www.stata.com)

Remarks are presented under the following headings:

Using bayesgraph

Examples

Trace plots

Autocorrelation plots

Histogram plots

Kernel density plots

Cumulative sum plots

Bivariate scatterplots

Diagnostic plots

Functions of model parameters

Using bayesgraph

`bayesgraph` requires specifying at least one parameter with all graph types except `matrix`, which requires at least two parameters. To request graphs for all parameters, use `_all`.

When multiple graphs are produced, they are automatically stored in memory with names `Graph_#` and will all appear on the screen. After you are done reviewing the graphs, you can type

```
. graph close Graph_*
```

to close these graphs or type

```
. graph drop Graph_*
```

to close the graphs and drop them from memory.

If you would like to see only one graph at a time, you can specify option `close` to close the Graph window when the next graph is displayed. You can also use option `sleep()` or option `wait` to pause between the subsequent graphs. The `sleep(#)` option causes each graph to pause for # seconds. The `wait` option causes `bayesgraph` to wait until a key is pressed before producing the next graph.

You can combine separate graphs into one by specifying one of `byparm()` or `combine()`. These options are not allowed with `diagnostics` or `matrix` graphs. The `byparm()` option produces more compact graphs, but it may not be feasible with many parameters or expressions and large sizes of MCMC samples.

With multiple graphs, you can control the look of each individual graph with `graph#opts()`. Options common to all graphs may be specified in `graphopts()` or passed directly to the command as with single graphs.

Examples

We demonstrate the `bayesgraph` command using an example of Bayesian normal linear regression applied to `auto.dta`. We model the `mpg` variable using a normal distribution with unknown mean and variance. Our Bayesian model thus has two parameters, `{mpg:_cons}` and `{var}`, for which we need to specify prior distributions. We consider fairly noninformative prior distributions for these parameters: $N(0, 1000)$ for the constant and inverse gamma with shape and scale of 0.1 for the variance. Because the specified prior distributions are independent and `semiconjugate` relative to the normal data distribution, we can use Gibbs sampling for both parameters instead of the default MH sampling. To illustrate, we will use Gibbs sampling for the variance and MH sampling (default) for the mean.

We use `bayesmh` to fit our model.

```
. use http://www.stata-press.com/data/r14/auto
(1978 Automobile Data)
. set seed 14
. bayesmh mpg, likelihood(normal({var}))
> prior({mpg:_cons}, normal(0,1000))
> prior({var}, igamma(0.1,0.1)) block({var}, gibbs)
Burn-in ...
Simulation ...
Model summary
```

```
Likelihood:
  mpg ~ normal({mpg:_cons},{var})
Priors:
  {mpg:_cons} ~ normal(0,1000)
  {var} ~ igamma(0.1,0.1)
```

```
Bayesian normal regression                MCMC iterations =    12,500
Metropolis-Hastings and Gibbs sampling    Burn-in          =     2,500
                                           MCMC sample size =   10,000
                                           Number of obs    =     74
                                           Acceptance rate  =    .7133
                                           Efficiency: min  =    .2331
                                           avg              =    .6166
                                           max              =     1
Log marginal likelihood = -242.1155
```

	Mean	Std. Dev.	MCSE	Median	Equal-tailed [95% Cred. Interval]	
mpg						
_cons	21.29231	.6648867	.013771	21.29419	19.94367	22.56746
var	34.2805	5.844213	.058442	33.6464	24.65882	47.5822

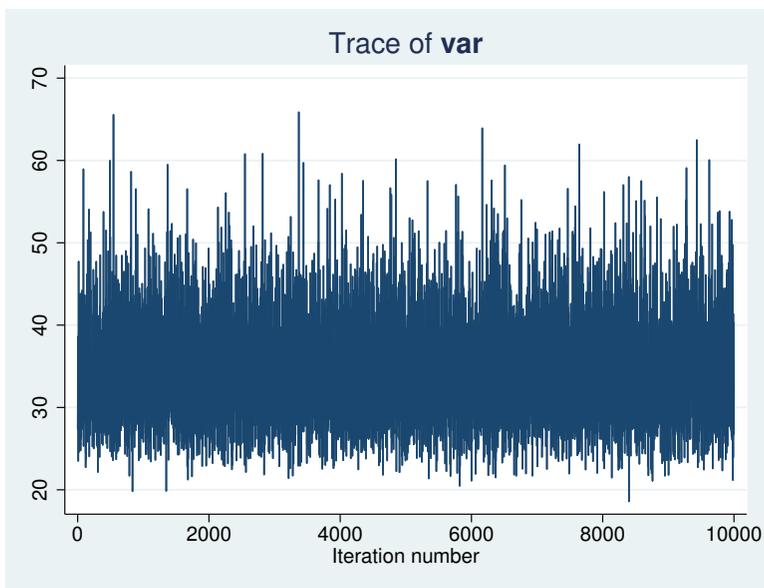
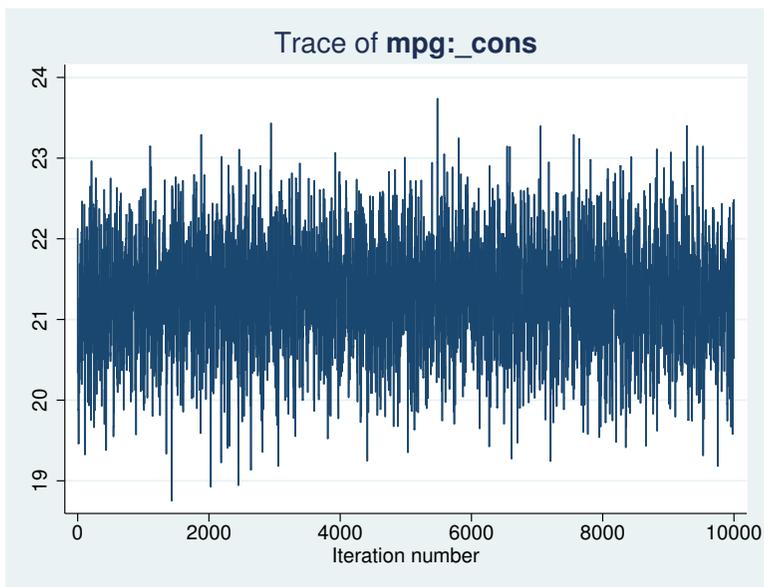
The MCMC simulation has a fairly high efficiency for the MH algorithm of 23% for the mean and an efficiency of 1 for the variance because of the Gibbs sampling. The output suggests no convergence problems. However, it is important to verify this and to also inspect various other graphical summaries of the parameters. This example demonstrates graphical summaries for a well-mixing MCMC chain that has converged and that generates samples from the posterior distribution of the model. For examples of poor-mixing MCMC chains, see *Convergence diagnostics in MCMC* in [\[BAYES\] intro](#).

Trace plots

We start with trace plots, which plot the values of the simulated parameters against the iteration number and connect consecutive values with a line. For a well-mixing parameter, the range of the parameter is traversed rapidly by the MCMC chain, which makes the drawn lines look almost vertical and dense. Sparseness and trends in the trace plot of a parameter suggest convergence problems.

Let's use `bayesgraph trace` to obtain trace plots for `{mpg:_cons}` and `{var}`. We specify `_all` to request both plots at once.

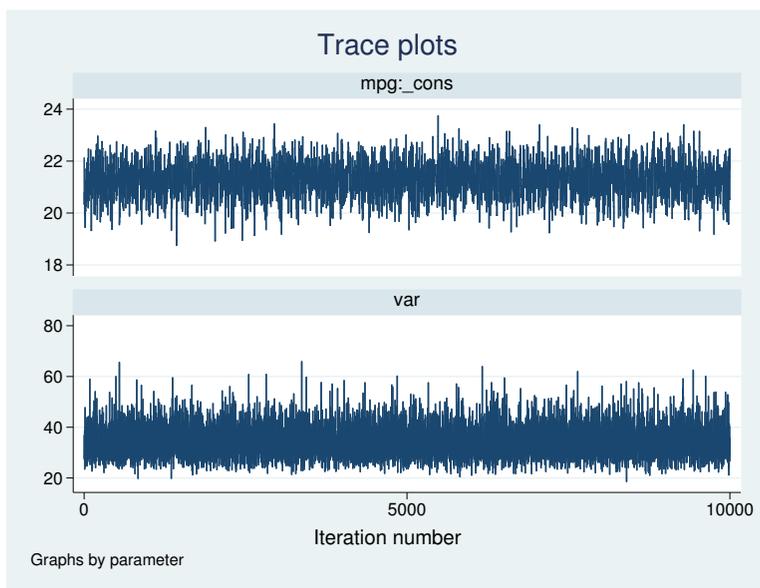
```
. bayesgraph trace _all
```



The mean parameter mixes very well and the variance parameter mixes perfectly.

Alternatively, we can use the `byparm()` option to plot results on one graph.

```
. bayesgraph trace _all, byparm
```



`bayesgraph trace` (as well as `bayesgraph cusum`) with option `byparm()` displays multiple plots in one column to accommodate an x axis with many iterations. You can specify `byparm(norowcoldefault)` to switch to the default behavior of options `rows()` and `cols()` documented in [G-3] *by_option*.

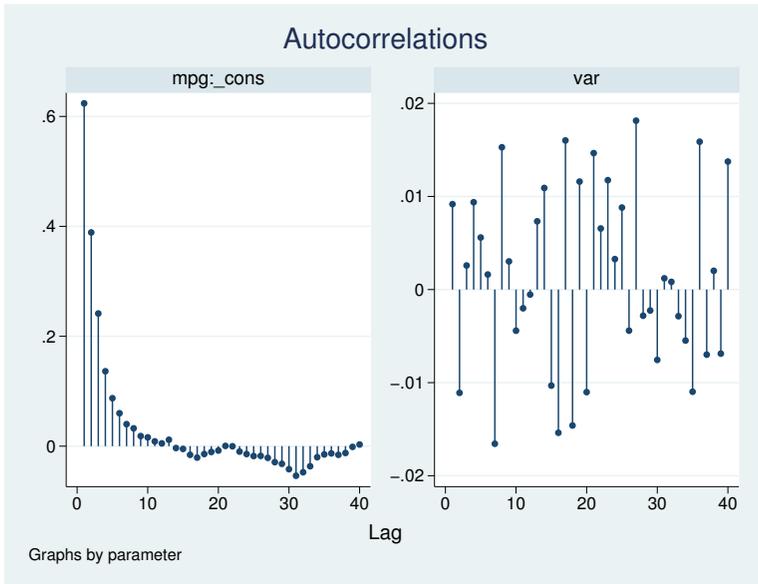
Autocorrelation plots

The second graphical summary we demonstrate is an autocorrelation plot. This plot shows the degree of autocorrelation in an MCMC sample for a range of lags, starting from lag 0. At lag 0, the plotted value corresponds to the sample variance of MCMC.

Autocorrelation is usually present in any MCMC sample. Typically, autocorrelation starts from some positive value for lag 0 and decreases toward 0 as the lag index increases. For a well-mixing MCMC chain, autocorrelation dies off fairly rapidly.

For example, autocorrelation for {mpg: _cons} becomes negligible after about lag 8 and is basically nonexistent for {var}.

```
. bayesgraph ac _all, byparm
```



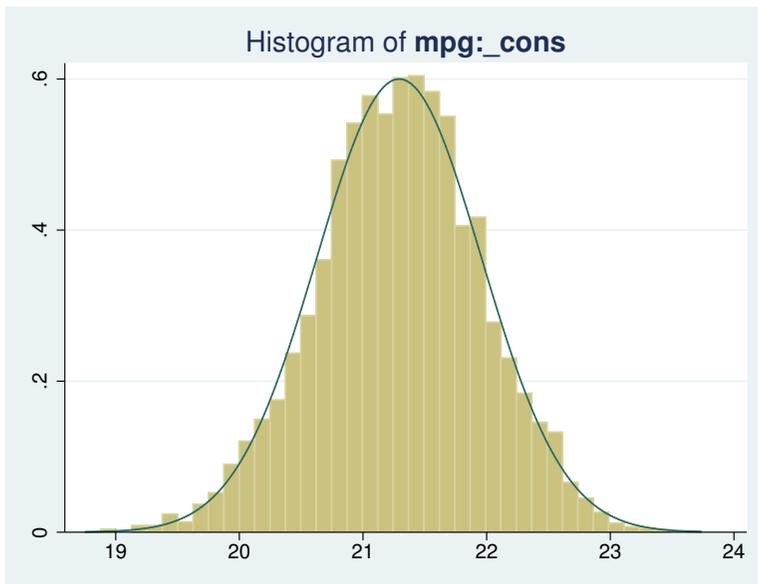
Autocorrelation lags are approximated by correlation times of parameters as reported by the `bayesstats ess` command; see [\[BAYES\] bayesstats ess](#) for details. Autocorrelation lags are also used to determine the batch size for the batch-means estimator of the MCMC standard errors; see [\[BAYES\] bayesstats summary](#).

Histogram plots

Graphical posterior summaries such as histograms and kernel density estimates provide useful additions to the various numerical statistics (see [\[BAYES\] bayesstats summary](#)) for summarizing MCMC output. It is always a good practice to inspect the histogram and kernel density estimates of the marginal posterior distributions of parameters to ensure that these empirical distributions behave as expected. These plots can be used to compare the empirical posterior and the specified prior distributions to visualize the impact of the data.

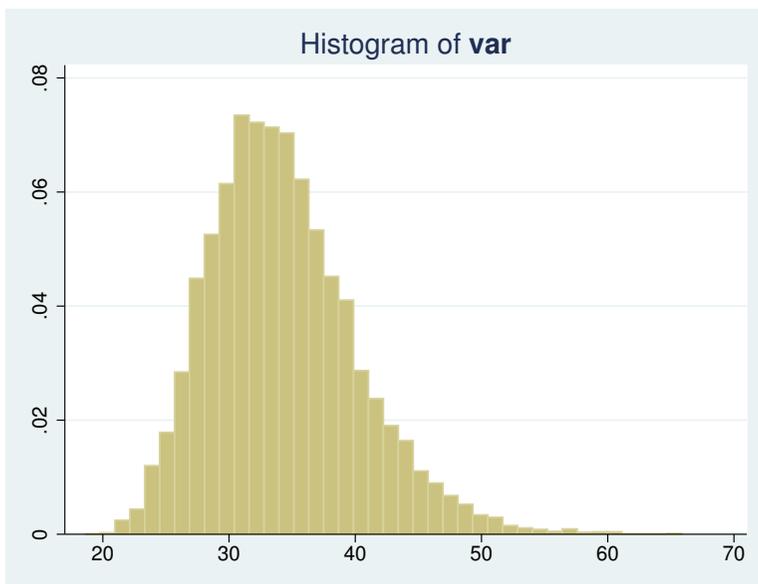
A histogram depicts the general shape of the marginal posterior distribution of a model parameter. Let's look at histograms of our parameters.

```
. bayesgraph histogram {mpg:_cons}, normal
```



The distribution of `{mpg:_cons}` is in good agreement with the normal distribution. This is not surprising, because the specified conjugate normal prior implies that the marginal posterior for `{mpg:_cons}` is a normal distribution. The unimodal histogram is also another confirmation that we have obtained a good simulation of the marginal posterior distribution of `{mpg:_cons}`.

```
. bayesgraph histogram {var}
```



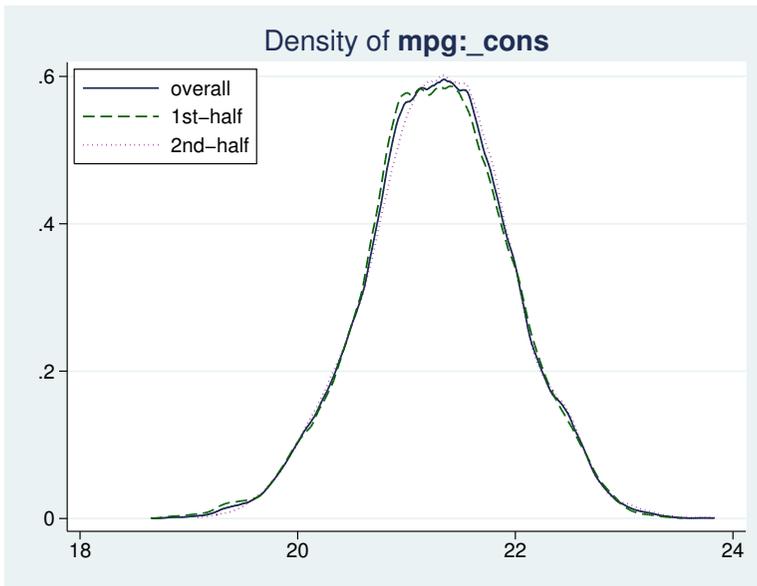
The histogram for `{var}` is also unimodal but is slightly skewed to the right. This is also in agreement with the specified prior because the marginal posterior for the variance is inverse gamma for the specified model.

Kernel density plots

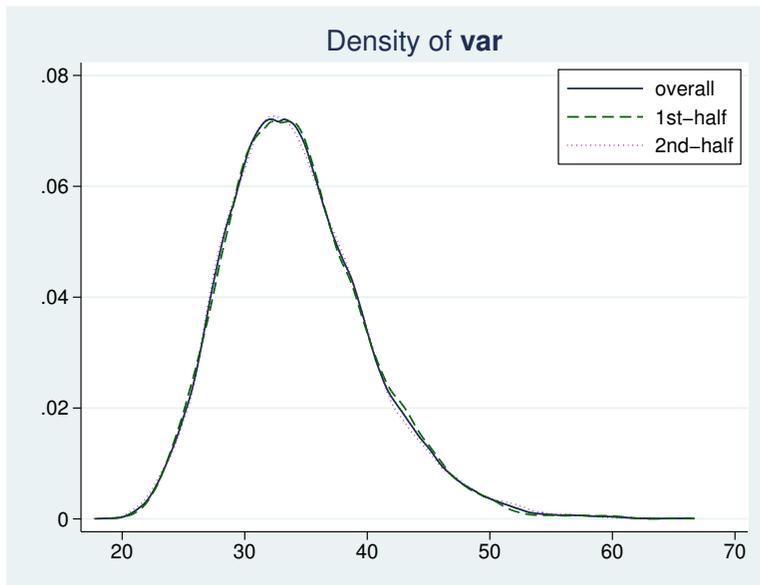
Kernel density plots provide alternative visualizations of the simulated marginal posterior distributions. They may be viewed as smoothed histograms. By default, the `bayesgraph kdensity` command shows three density curves: an overall density of the entire MCMC sample, the first-half density obtained using the first half of the MCMC sample, and the second-half density obtained using the second half of the MCMC sample. If the chain has converged and mixes well, we expect these three density curves to be close to each other. Large discrepancies between the first-half curve and the second-half curve suggest convergence problems.

Let's look at kernel density plots for our two parameters.

```
. bayesgraph kdensity {mpg:_cons}
```



```
. bayesgraph kdensity {var}
```



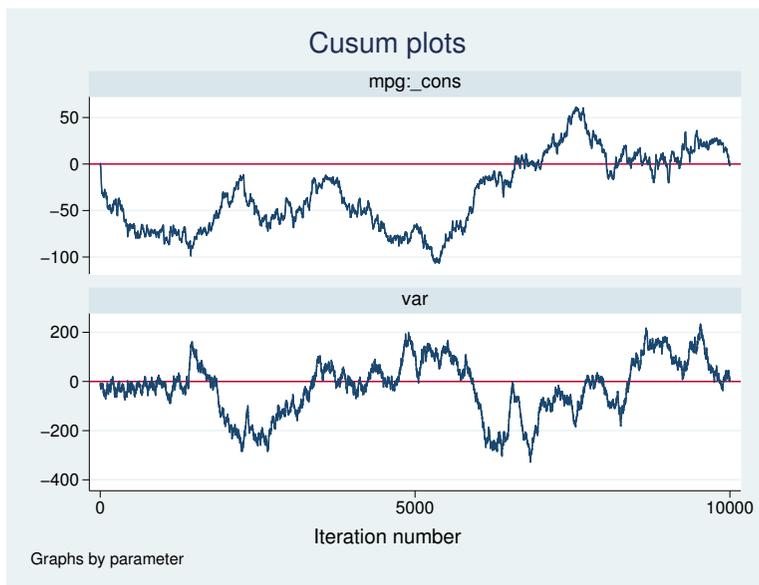
Kernel density plots for `{mpg:_cons}` and `{var}` are similar in shape to the histograms' plots from the previous section. All three density curves are close to each other for both parameters.

Cumulative sum plots

Cumulative sum (cusum) plots are useful graphical summaries for detecting persistent trends in MCMC chains. All cusum plots start and end at 0 and may or may not cross the x axis. There is great variability in the looks of cusum plots, which make them difficult to interpret sometimes. Typically, if the cusum line never crosses the x axis, this may indicate a problem. See, for example, [Convergence diagnostics of MCMC](#) in [\[BAYES\] intro](#) for a cusum plot demonstrating convergence problems.

By inspecting a cusum plot, we may detect an early drift in the simulated sample because of an insufficient burn-in period. In cases of pronounced persistent trends, the cusum curve may stay either in the positive or in the negative y plane. For a well-mixing parameter, the cusum curve typically crosses the x axis several times. This is the case for the cusum plots of `{mpg:_cons}` and `{var}`.

```
. bayesgraph cusum _all, byparm
```



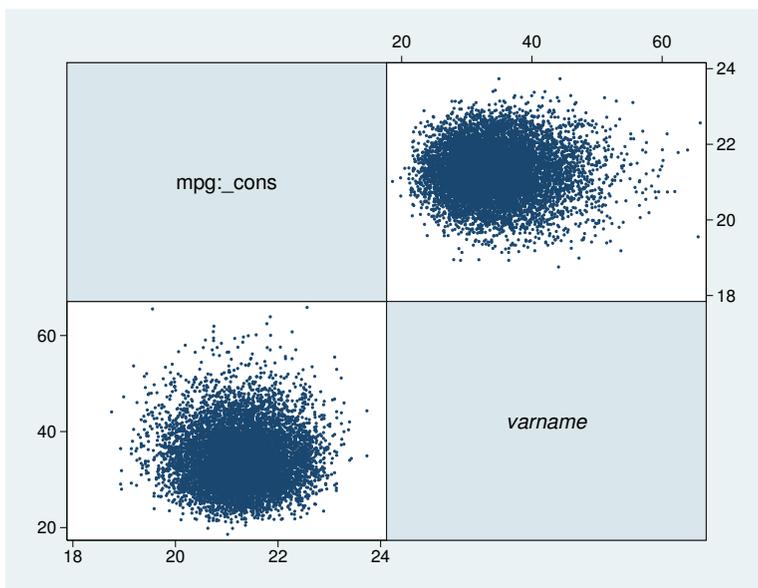
Bivariate scatterplots

The `bayesgraph matrix` command draws bivariate scatterplots of model parameters based on MCMC samples. A bivariate scatterplot represents a joint sample posterior distribution for pairs of parameters. It may reveal correlation between parameters and characterize a general shape of a multivariate posterior distribution. For example, bivariate scatterplots are useful for detecting multimodal posterior distributions.

Typically, scatterplots depict clouds of points. Sparseness and irregularities in the scatterplots can be strong indications of nonconvergence of an MCMC. For a well-mixing chain, the scatterplots have an ellipsoidal form with an increasing concentration around the posterior mode.

This scatterplot of {mpg:_cons} and {var} is an example of a well-behaved scatterplot.

```
. bayesgraph matrix {mpg:_cons} {var}
```

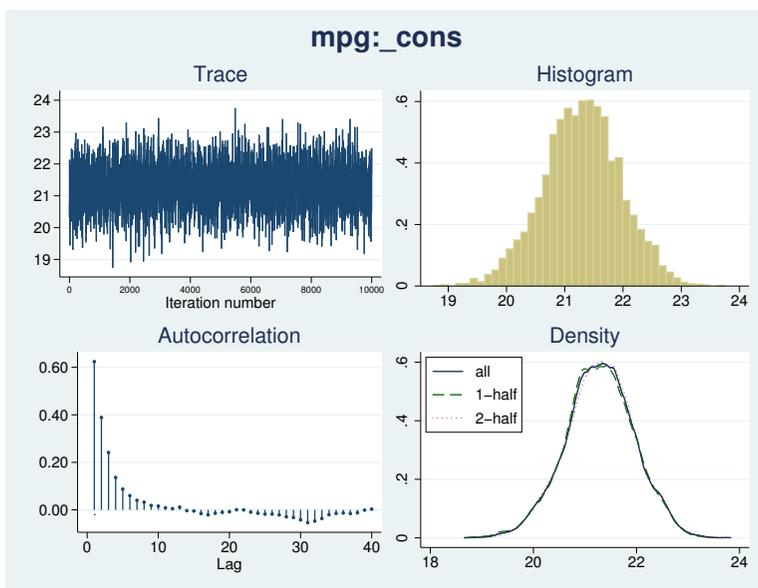


Diagnostic plots

Finally, we demonstrate the `bayesgraph diagnostics` command, which combines the trace, histogram, autocorrelation, and kernel density plots compactly on one graph. We already discussed the individual plots in the previous sections. Diagnostic plots are convenient for inspecting the overall behavior of a particular model parameter. We recommend that diagnostic plots for all parameters be inspected routinely as a part of the convergence-checking process.

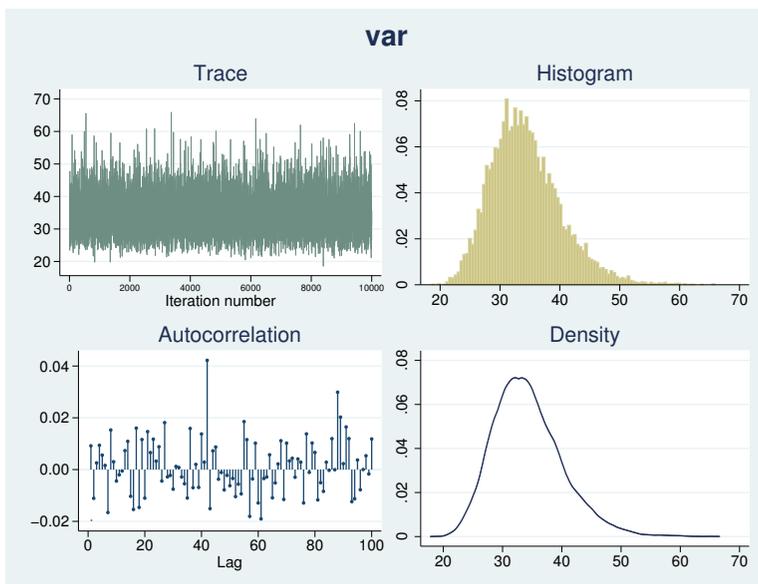
Let's obtain the diagnostic plot for {mpg:_cons}.

```
. bayesgraph diagnostics {mpg:_cons}
```



In the diagnostics plot for {var}, let's also demonstrate the use of several options of the depicted plots.

```
. bayesgraph diagnostics {var}, traceopts(lwidth(0.2) lcolor(teal))
> acopts(lag(100)) histoopts(bin(100)) kdensopts(show(none))
```

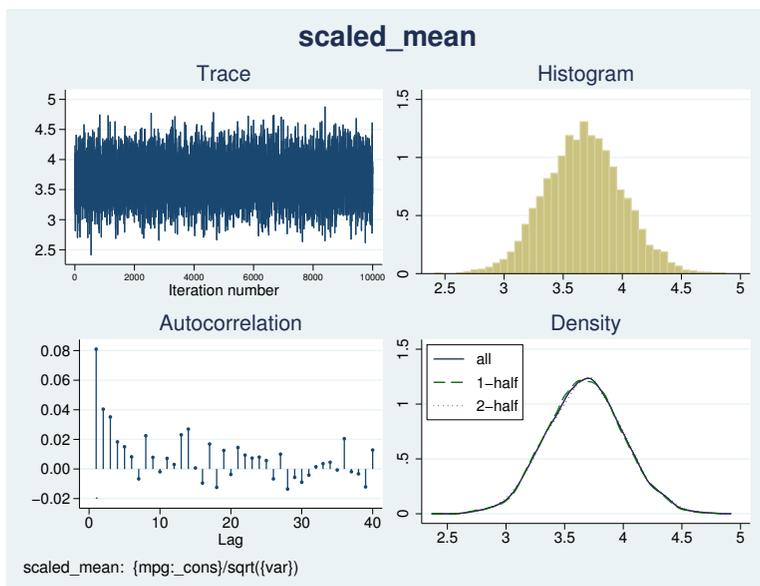


In the above, we changed the width and color of the trace line, the maximum lag for calculating the autocorrelation, the number of bins for the histogram, and requested that the two subsample kernel densities not be shown on the kernel density plot.

Functions of model parameters

All `bayesgraph` subcommands can provide graphical summaries of functions of model parameters. Below we apply `bayesgraph` diagnostics to the expression `{mpg:_cons}/sqrt({var})`, which we label as `scaled_mean`.

```
. bayesgraph diagnostics (scaled_mean: {mpg:_cons}/sqrt({var}))
```



If you detect convergence problems in a function of parameters, you must inspect every parameter used in the expression individually. In fact, we recommend that you inspect all model parameters before you proceed with any postestimation analysis.

Methods and formulas

Let θ be a scalar model parameter and $\{\theta_t\}_{t=1}^T$ be an MCMC sample of size T drawn from the marginal posterior distribution of θ .

The trace plot of θ plots θ_t against t with connecting lines for $t = 1, \dots, T$.

The autocorrelation plot of θ shows the autocorrelation in the $\{\theta_t\}_{t=1}^T$ sample for lags from 0 to the `lag(#)` option of the `ac` command.

The histogram and kernel density plots of θ are drawn using the `histogram` and `kdensity` commands.

Yu and Mykland (1998) proposed a graphical procedure for assessing the convergence of individual parameters based on cumulative sums, also known as a cusum plot. The cusum plot for θ plots S_t against t for $t = 1, \dots, T$ and connects the successive points. S_t is the cumulative sum at time t :

$$S_t = \sum_{k=1}^t (\theta_k - \hat{\theta}), \quad \hat{\theta} = \frac{1}{T} \sum_{k=1}^T \theta_k$$

and $S_0 = 0$.

The scatterplot of two model parameters θ^1 and θ^2 plots points (θ_t^1, θ_t^2) for $t = 1, \dots, T$.

Reference

Yu, B., and P. Mykland. 1998. Looking at Markov samplers through cusum path plots: A simple diagnostic idea. *Statistics and Computing* 8: 275–286.

Also see

[BAYES] [bayesmh](#) — Bayesian regression using Metropolis–Hastings algorithm

[BAYES] [bayesmh postestimation](#) — Postestimation tools for bayesmh

[BAYES] [bayesstats ess](#) — Effective sample sizes and related statistics

[BAYES] [bayesstats summary](#) — Bayesian summary statistics

[G-2] [graph matrix](#) — Matrix graphs

[G-2] [graph twoway kdensity](#) — Kernel density plots

[R] [histogram](#) — Histograms for continuous and categorical variables

[R] [kdensity](#) — Univariate kernel density estimation

[TS] [corrgram](#) — Tabulate and graph autocorrelations

[TS] [tline](#) — Plot time-series data