bmagraph msize — Model-size distribution plots after BMA regression

Description

`bmagraph msize` provides a graphical summary for the posterior and prior model-size distributions after the `bmaregress` command.

Quick start

Plot posterior and prior model-size distributions after fitting a Bayesian model averaging (BMA) linear regression

```
bmagraph msize
```

Same as above, but plot the posterior model-size distribution only

```
bmagraph msize, noprior
```

Menu

Statistics > Bayesian model averaging > Model-size distributions

Syntax

`bmagraph msize` [, *options*]

options	Description
Main	
`constant`	include constant term in model-size computations; default is no constant
`noprior`	suppress prior model-size distribution shown by default
`priorlineopts(`*cline_options*`)`	affect rendition of prior line
`anlineopts(`*cline_options*`)`	affect rendition of analytical posterior line; ignored with random g
`freqlineopts(`*cline_options*`)`	affect rendition of frequency posterior line; ignored with model enumeration
`postlineopts(`*cline_options*`)`	affect rendition of all posterior lines
Line options	
cline_options	affect rendition of all plotted lines
Y axis, X axis, Titles, Legend, Overall	
twoway_options	any options other than `by()` documented in [G-3] ***twoway_options***

Options

Main

`constant` specifies that the constant term be included in model-size computations. By default, the constant term is not included.

`noprior` specifies that the prior model-size distribution not be shown on the plot. It is shown by default.

`priorlineopts(`*cline_options*`)` affects the rendition of the prior line; see [G-3] ***cline_options***.

`anlineopts(`*cline_options*`)` affects the rendition of the analytical posterior line; see [G-3] ***cline_options***. This option is ignored for BMA models with a random g.

`freqlineopts(`*cline_options*`)` affects the rendition of the frequency posterior line; see [G-3] ***cline_options***. The frequency model-size distribution is plotted whenever sampling is used. This option is ignored with model enumeration.

`postlineopts(`*cline_options*`)` affects the rendition of the analytical and frequency posterior lines; see [G-3] ***cline_options***.

Line options

cline_options affects the rendition of all plotted lines; see [G-3] ***cline_options***.

Y axis, X axis, Titles, Legend, Overall

twoway_options are any of the options documented in [G-3] ***twoway_options***, excluding `by()`. These include options for titling the graph (see [G-3] ***title_options***) and for saving the graph to disk (see [G-3] ***saving_option***).

Remarks and examples

stata.com

See *Remarks and examples* in [BMA] **bmastats msize** for a general discussion of a BMA model size.

A model-size distribution is used to explore model complexity. The prior model-size distribution describes our a priori assumption about model size. The posterior model-size distribution describes the effect of the data on the BMA model. For this purpose, we compare the posterior and prior model-size distributions. If the posterior is skewed to the left with respect to the prior, then the data favor smaller models than assumed by the prior. If the posterior is skewed to the right, then the data favor larger models than assumed by the prior.

For the prior model-size distribution, `bmagraph msize` always uses analytical computation. In the case of sampling, it is conditional on the visited models. For the posterior model-size distribution, it plots the analytical distribution for a fixed g and the MCMC frequency-based or simply frequency distribution for a random g and whenever sampling is used.

▷ Example 1: Model-size distributions for BMA models using enumeration

We consider the `performance` dataset (Chatterjee and Hadi 2012, sec. 3.3) analyzed in example 1 of [BMA] **bmastats msize**. In that example, we explored only a few summaries of the model-size distributions. Here we describe the entire distributions.

The `rating` variable is regressed on all predictors from `complaints` to `advance`. By default, because of the small number of predictors, six, the model space is explored fully by using enumeration.

```
. use https://www.stata-press.com/data/r18/performance
(Data on employee satisfaction with supervisor)

. bmaregress rating complaints-advance

Enumerating models ...
Computing model probabilities ...

Bayesian model averaging                    No. of obs          =      30
Linear regression                           No. of predictors   =       6
Model enumeration                                       Groups  =       6
                                                        Always  =       0
Priors:                                     No. of models       =      64
  Models: Beta-binomial(1, 1)                   For CPMP >= .9  =      10
   Cons.: Noninformative                    Mean model size     =   1.699
   Coef.: Zellner's g
       g: Benchmark, g = 36                 Shrinkage, g/(1+g)  = 0.9730
  sigma2: Noninformative                    Mean sigma2         =  52.302

-----------------------------------------------------------------------------
      rating |     Mean    Std. dev.                       Group         PIP
-------------+---------------------------------------------------------------
  complaints |  .7052859   .1224289                           1       .99973
    learning |  .0603014   .1285281                           3       .25249
     advance | -.0167921    .073883                           6       .13148
  privileges | -.0074174   .0488635                           2       .10998
      raises |  .0070789   .0670475                           4       .10642
    critical |  .0009713   .0437848                           5      .098534
-------------+---------------------------------------------------------------
Always       |
       _cons |   14.8472   7.874219                           0            1
-----------------------------------------------------------------------------
Note: Coefficient posterior means and std. dev. estimated from 64 models.
Note: Default priors are used for models and parameter g.
```

There is a total of $2^6 = 64$ models in the fully explored model space. Let's use `bmagraph msize` to draw the posterior and prior model-size distributions.

```
. bmagraph msize
note: frequency posterior model-size distribution not available.
```

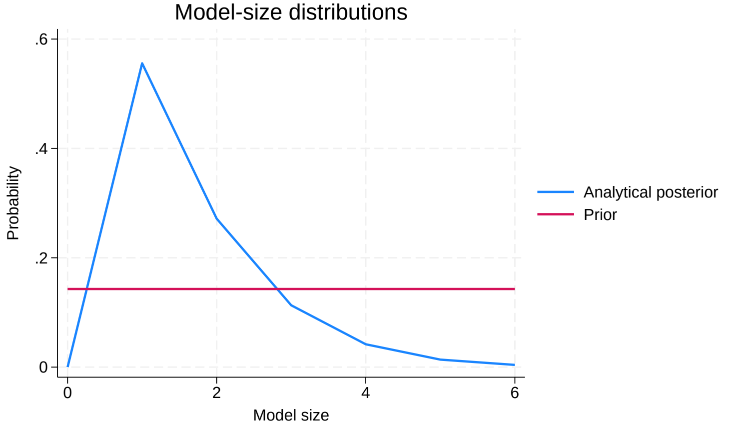

The reported model size does not include the constant, so its range is between 0 and 6. You may include the constant by specifying the `constant` option. By default, an uninformative uniform prior is assumed for the model size. The posterior model-size distribution is skewed to the left. Its mode is 1, so the posterior favors smaller models.

A frequency-based posterior estimate of the model-size distribution is not available in this example because there is no MCMC sample with model enumeration.

◁

▷ Example 2: Model-size distributions for BMA models using MCMC model composition (MC3) sampling

Continuing with example 1, we fit the same BMA model, but this time we use the MC3 sampling algorithm by specifying the `sampling` option. We also specify the `rseed()` option for reproducibility.

```
. bmaregress rating complaints-advance, sampling rseed(18)
Burn-in ...
Simulation ...
Computing model probabilities ...
Bayesian model averaging                        No. of obs         =      30
Linear regression                               No. of predictors  =       6
MC3 sampling                                                Groups =       6
                                                            Always =       0
                                                No. of models      =      32
                                                    For CPMP >= .9 =      10
Priors:                                         Mean model size    =   1.699
  Models: Beta-binomial(1, 1)                   Burn-in            =   2,500
   Cons.: Noninformative                        MCMC sample size   =  10,000
   Coef.: Zellner's g                           Acceptance rate    =  0.2417
       g: Benchmark, g = 36                     Shrinkage, g/(1+g) =  0.9730
  sigma2: Noninformative                        Mean sigma2        =  52.292
Sampling correlation = 0.9990
```

rating	Mean	Std. dev.	Group	PIP
complaints	.705479	.1218881	1	1
learning	.0601919	.1282869	3	.25234
advance	-.0167514	.0737415	6	.13141
privileges	-.0074265	.048844	2	.10996
raises	.0069949	.0666406	4	.10629
critical	.0009699	.0437742	5	.098526
Always				
_cons	14.84478	7.871046	0	1

```
Note: Coefficient posterior means and std. dev. estimated from 32 models.
Note: Default priors are used for models and parameter g.
```

Instead of enumerating models (fully exploring the space), `bmaregress` explored only half the model space. It visited 32 out of the total 64 models. We inspect the effect of this on the prior and posterior model-size distributions.

```
. bmagraph msize
```

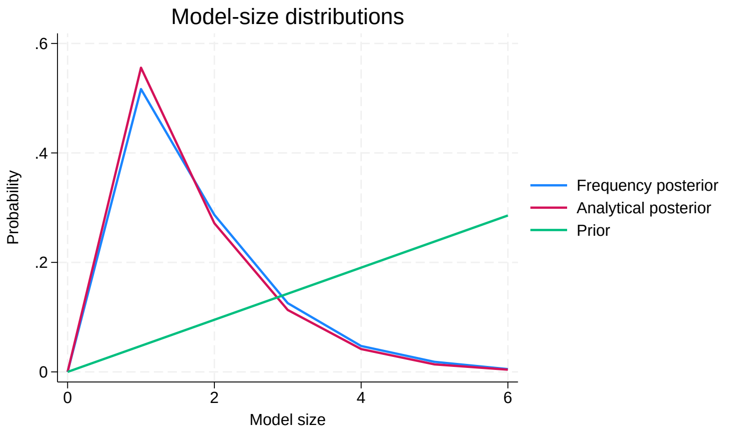

Although we used the same model prior as in example 1, the prior model-size distribution looks different. This is because our explored model space now contains 32 models instead of all 64, and the prior model-size distribution is conditional on the visited models.

The model of size 1, the one that includes `complaints`, has the highest posterior probability of about 0.56.

With a fixed g when we fit a BMA model using MC3 sampling, in addition to the analytical model-size distribution, the frequency posterior model-size distribution is available. Provided that the model-space sampling converges, the analytical and frequency distributions should be close. In our example, the analytical and frequency model-size distributions are nearly identical.

◁

▷ Example 3: Model-size distributions for BMA models with random g

Both example 1 and example 2 used a fixed g. Let's explore the case of a random g. (An in-depth coverage of the effects of the g-prior on model complexity can be found in, for example, Ley and Steel [2012].)

To demonstrate, we will use a `robust` prior for g.

```
. bmaregress rating complaints-advance, gprior(robust) rseed(18)
Burn-in ...
Simulation ...
Computing model probabilities ...

Bayesian model averaging                         No. of obs         =      30
Linear regression                                No. of predictors =       6
MC3 and adaptive MH sampling                                Groups =       6
                                                            Always =       0
                                                 No. of models      =      34
                                                    For CPMP >= .9 =      12
Priors:                                          Mean model size    =   1.734
  Models: Beta-binomial(1, 1)                    Burn-in            =   2,500
   Cons.: Noninformative                         MCMC sample size   =  10,000
   Coef.: Zellner's g                            Acceptance rate    =  0.4232
       g: Robust
  sigma2: Noninformative                         Mean sigma2        =  53.095

Sampling correlation = 0.9994

------------------------------------------------------------------------------
      rating |      Mean   Std. dev.                          Group        PIP
-------------+----------------------------------------------------------------
  complaints |  .7000463   .1273543                               1      .9998
    learning |  .0594904   .1286095                               3        .25
     advance | -.0192712   .0797935                               6      .1503
      raises |  .0079416   .0727859                               4      .1201
  privileges | -.0072591   .0487009                               2      .1069
    critical |  .0014397   .0466476                               5      .1067
-------------+----------------------------------------------------------------
Always       |
       _cons |  15.24911   7.988166                               0          1
------------------------------------------------------------------------------
Note: Coefficient posterior means and std. dev. estimated from 34 models.
Note: Default prior is used for models.

------------------------------------------------------------------------------
             |                                                Equal-tailed
             |      Mean   Std. dev.      MCSE     Median  [95% cred. interval]
-------------+----------------------------------------------------------------
           g |   152.668   1968.132    43.5265   33.81024   8.205076   610.6026
   Shrinkage |  .9656427   .0276071    .001234   .9712728   .8913639   .9983649
------------------------------------------------------------------------------
```

`bmaregress` now uses MC3 sampling for the models and adaptive Metropolis–Hastings sampling for g.

```
. bmagraph msize
note: analytical posterior model-size distribution not available.
```

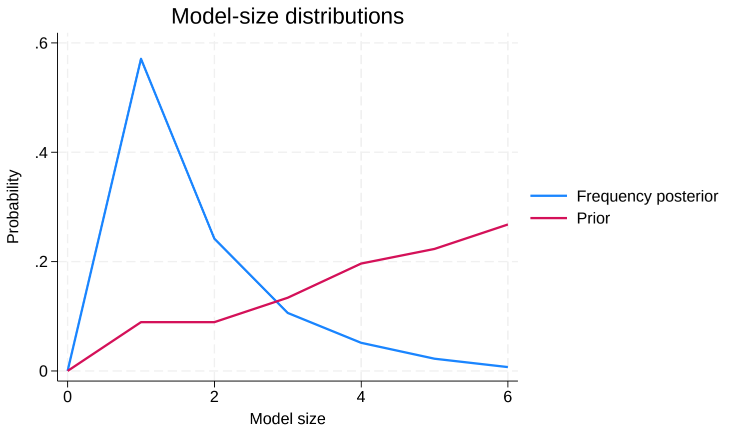

◁

The analytical posterior model-size distribution is not available with a random g. The frequency posterior model-size distribution is similar to that in example 2 for fixed $g = 36$. Particularly, the null model was not visited by the MC3 sampler, and models of size 1 and 2 have the highest posterior probabilities (both above 0.2), but the mode of the posterior distribution here is 1.

Methods and formulas

See *Methods and formulas* in [BMA] **bmastats msize**.

References

Chatterjee, S., and A. S. Hadi. 2012. *Regression Analysis by Example*. 5th ed. New York: Wiley.

Ley, E., and M. F. J. Steel. 2012. Mixtures of g-priors for Bayesian model averaging with economic applications. *Journal of Econometrics* 171: 251–266. https://doi.org/10.1016/j.jeconom.2012.06.009.

Also see

[BMA] **bmastats msize** — Model-size summary after BMA regression

[BMA] **bmagraph** — Graphical summary for models and predictors after BMA regression

[BMA] **bmaregress** — Bayesian model averaging for linear regression

[BMA] **BMA postestimation** — Postestimation tools for Bayesian model averaging

[BMA] **Glossary**