Bayesian Econometrics in Stata 17

David Schenck

Senior Econometrician
Stata

2021 US Stata Conference
August 6, 2021

Bayesian econometrics in Stata

- Stata 17 introduces Bayesian estimation of a variety of time–series and panel–data econometric models
  - Multivariate time-series:
    - `bayes: var`
    - `bayes: dsge`
    - `bayes: dsgenl`
    - `bayesirf`
    - `bayesfcast`
  - Panel data:
    - `bayes: xtreg`
    - `bayes: xtlogit`
    - `bayes: xtprobit`
    - `bayes: xtologit`
    - `bayes: xtoprobit`
    - `bayes: xtmlogit`
    - `bayes: xtpoisson`
    - `bayes: xtnbreg`

`bayes: var`

The vector autoregression model

- A VAR expresses a collection of variables as functions of their lags

$$\mathbf{y}_t = \mathbf{a} + \mathbf{A}_1 \mathbf{y}_{t-1} + \cdots + \mathbf{A}_p \mathbf{y}_{t-p} + \mathbf{C}\mathbf{x}_t + \mathbf{u}_t \qquad \mathbf{u}_t \sim N(\mathbf{0}, \boldsymbol{\Sigma})$$

- $\mathbf{y}_t$ is a vector of k variables
- $\mathbf{u}_t$ is a vector of k disturbance terms with $k \times k$ covariance matrix $\boldsymbol{\Sigma}$
- $\mathbf{A}_i$ is a $k \times k$ matrix of parameters for $i = 1, 2, \ldots, p$
- can include exogenous variables $\mathbf{x}_t$ with coefficients $\mathbf{C}$
- Structure is minimal: choice of k variables, p lags

VAR estimation

- Flexible setup with minimal structure
- But: many parameters to estimate ($k^2 p$ slope coefficients, k constant terms, $k(k+1)/2$ elements of $\boldsymbol{\Sigma}$)
- The large number of parameters to be estimated can lead to imprecise estimates

US macro data

```
. webuse usmacro
(Federal Reserve Economic Data - St. Louis Fed)

. describe

Contains data from https://www.stata-press.com/data/r17/usmacro.dta
 Observations:           226                  Federal Reserve Economic Data -
                                                St. Louis Fed
    Variables:             4                  4 Dec 2020 12:39
-------------------------------------------------------------------------------
Variable      Storage   Display    Value
    name         type    format    label      Variable label
-------------------------------------------------------------------------------
fedfunds        double  %10.0g                Federal funds rate
date            int     %tq                   Date (quarters)
inflation       float   %9.0g                 Annual rate of inflation
ogap            float   %9.0g                 GDP gap
-------------------------------------------------------------------------------
Sorted by: date
```

US macro data

```
. tsline inflation ogap fedfunds
```

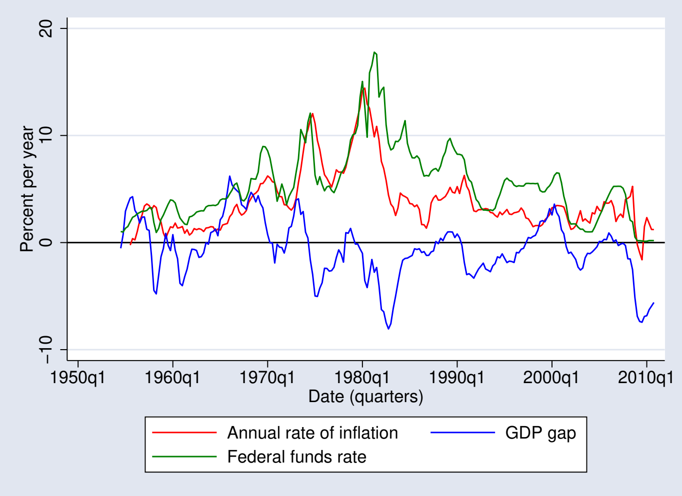

A simple VAR

```
. var ogap inflation fedfunds if date <= tq(2005q1), lags(1/4)

Vector autoregression

Sample: 1956q3 thru 2005q1                      Number of obs     =        195
Log likelihood =  -578.7363                     AIC               =   6.335757
FPE            =   .1133494                     HQIC              =   6.600797
Det(Sigma_ml)  =   .0759353                     SBIC              =   6.990357

Equation           Parms      RMSE     R-sq      chi2     P>chi2
----------------------------------------------------------------

ogap                 13     .777444   0.9138   2066.528   0.0000
inflation            13     .507601   0.9715    6637.26   0.0000
fedfunds             13     .830675   0.9420   3166.504   0.0000
----------------------------------------------------------------
```

(output continues on next slide)

A simple VAR

	Coefficient	Std. err.	z	P>|z|	[95% conf.	interval]
ogap						
ogap						
L1.	1.138529	.0731992	15.55	0.000	.9950609	1.281996
L2.	-.0570314	.1075964	-0.53	0.596	-.2679165	.1538536
L3.	-.1926199	.1041447	-1.85	0.064	-.3967398	.0115
L4.	.0220416	.0726475	0.30	0.762	-.1203449	.1644281
inflation						
L1.	-.0599264	.111144	-0.54	0.590	-.2777647	.1579119
L2.	.1007764	.1762956	0.57	0.568	-.2447566	.4463094
L3.	-.1865833	.1776786	-1.05	0.294	-.5348269	.1616603
L4.	.1257294	.1099893	1.14	0.253	-.0898457	.3413044
fedfunds						
L1.	.1003059	.0717657	1.40	0.162	-.0403524	.2409641
L2.	-.3828584	.0991487	-3.86	0.000	-.5771862	-.1885305
L3.	.3411735	.100363	3.40	0.001	.1444657	.5378813
L4.	-.1045785	.0724888	-1.44	0.149	-.246654	.037497
_cons	.3101242	.119271	2.60	0.009	.0763575	.543891

(output continues on next slide)

A simple VAR

inflation						
ogap						
L1.	.0525665	.0477925	1.10	0.271	-.0411051	.1462381
L2.	-.0234909	.0702508	-0.33	0.738	-.16118	.1141982
L3.	.0305625	.0679972	0.45	0.653	-.1027096	.1638346
L4.	-.0364685	.0474323	-0.77	0.442	-.1294342	.0564971
inflation						
L1.	1.205197	.0725671	16.61	0.000	1.062968	1.347426
L2.	-.2206968	.1151053	-1.92	0.055	-.446299	.0049054
L3.	.1078881	.1160083	0.93	0.352	-.119484	.3352601
L4.	-.1358476	.0718132	-1.89	0.059	-.2765988	.0049037
fedfunds						
L1.	.1866136	.0468566	3.98	0.000	.0947763	.2784509
L2.	-.1912048	.0647353	-2.95	0.003	-.3180836	-.0643261
L3.	.1339269	.0655281	2.04	0.041	.0054942	.2623595
L4.	-.1261127	.0473287	-2.66	0.008	-.2188753	-.0333501
_cons	.1772255	.0778733	2.28	0.023	.0245967	.3298544

(output continues on next slide)

A simple VAR

fedfunds						
ogap						
L1.	.3937191	.0782111	5.03	0.000	.2404281	.5470101
L2.	-.1511135	.1149635	-1.31	0.189	-.3764379	.0742109
L3.	-.1325565	.1112755	-1.19	0.234	-.3506525	.0855396
L4.	-.0665638	.0776217	-0.86	0.391	-.2186995	.0855719
inflation						
L1.	-.0239604	.1187541	-0.20	0.840	-.2567142	.2087933
L2.	.4286425	.1883666	2.28	0.023	.0594508	.7978342
L3.	-.2639582	.1898443	-1.39	0.164	-.6360461	.1081297
L4.	-.0248893	.1175203	-0.21	0.832	-.2552248	.2054462
fedfunds						
L1.	1.082223	.0766795	14.11	0.000	.9319335	1.232512
L2.	-.5203612	.1059374	-4.91	0.000	-.7279948	-.3127277
L3.	.455501	.1072348	4.25	0.000	.2453246	.6656775
L4.	-.1083045	.0774521	-1.40	0.162	-.2601079	.0434989
_cons	.089251	.1274375	0.70	0.484	-.1605218	.3390238

(output completes)

The Bayesian VAR

- Likelihood model: observables $\mathbf{y}_t$ follow

$$\mathbf{y}_t = \mathbf{a} + \mathbf{A}_1\mathbf{y}_{t-1} + \cdots + \mathbf{A}_p\mathbf{y}_{t-p} + \mathbf{C}\mathbf{x}_t + \mathbf{u}_t, \qquad \mathbf{u}_t \sim N(\mathbf{0}, \boldsymbol{\Sigma})$$

- Prior for coefficients $\boldsymbol{\beta} = vec(\mathbf{A}_1, \ldots, \mathbf{A}_p, \mathbf{C})$ is multivariate normal

$$\boldsymbol{\beta} \sim N(\boldsymbol{\beta}_0, \boldsymbol{\Omega})$$

- Prior for covariance matrix $\boldsymbol{\Sigma}$ is either inverse Wishart or Jeffreys

Bayesian VAR priors I

- Look at a two-variable VAR with two lags for simplicity:

$$y_t = a_{11} y_{t-1} + a_{12} p_{t-1} + b_{11} y_{t-2} + b_{12} p_{t-2} + u_{1t}$$
$$p_t = a_{21} y_{t-1} + a_{22} p_{t-1} + b_{21} y_{t-2} + b_{22} p_{t-2} + u_{2t}$$

- Three types of slope coefficients:
  1. First autoregressive lag (red)
  2. Other autoregressive lags (blue)
  3. Cross-lags (black)
- Priors are specified for each type of coefficient

Bayesian VAR priors II

- Priors are normally distributed
- Prior means:
  - First autoregressive lag: prior mean 1
  - Further AR lags and all cross-lags: prior mean 0
- Prior variances:

$$\text{Autoregressive lags} = \left(\frac{\lambda_1}{l^{\lambda_3}}\right)^2 \qquad \text{Cross lags} = \frac{\sigma_i^2}{\sigma_j^2}\left(\frac{\lambda_1 \lambda_2}{l^{\lambda_3}}\right)^2$$

- Interpretation:
  - λ_1: autoregressive lag tightness (default: 0.1)
  - λ_2: cross-lag tightness (default: 0.5)
  - λ_3: lag attenuation (default: 1)

Bayesian VAR priors III

- The upshot: random walk prior, with variances that are tighter around 0 as the lag length increases

$$y_t = a_{11} y_{t-1} + a_{12} p_{t-1} + b_{11} y_{t-2} + b_{12} p_{t-2} + u_{1t}$$
$$p_t = a_{21} y_{t-1} + a_{22} p_{t-1} + b_{21} y_{t-2} + b_{22} p_{t-2} + u_{2t}$$

Bayesian VAR priors III

- The upshot: random walk prior, with variances that are tighter around 0 as the lag length increases

$$y_t = a_{11} y_{t-1} + a_{12} p_{t-1} + b_{11} y_{t-2} + b_{12} p_{t-2} + u_{1t}$$

$$p_t = a_{21} y_{t-1} + a_{22} p_{t-1} + b_{21} y_{t-2} + b_{22} p_{t-2} + u_{2t}$$

bayes: var

```
. bayes, rseed(17): var ogap inflation fedfunds if date <= tq(2005q1), lag(1/4)

Burn-in ...
Simulation ...

Model summary
------------------------------------------------------------------------------
Likelihood:
       ogap
  inflation
   fedfunds ~ mvnormal(3,xb_ogap,xb_inflation,xb_fedfunds,{Sigma,m})

Priors:
            {ogap:L(1 2 3 4).ogap}                                          (1)
       {ogap:L(1 2 3 4).inflation}                                          (1)
        {ogap:L(1 2 3 4).fedfunds}                                          (1)
                      {ogap:_cons}                                          (1)
       {inflation:L(1 2 3 4).ogap}                                          (2)
  {inflation:L(1 2 3 4).inflation}                                          (2)
   {inflation:L(1 2 3 4).fedfunds}                                          (2)
                 {inflation:_cons}                                          (2)
        {fedfunds:L(1 2 3 4).ogap}                                          (3)
   {fedfunds:L(1 2 3 4).inflation}                                          (3)
    {fedfunds:L(1 2 3 4).fedfunds}                                          (3)
                  {fedfunds:_cons} ~ varconjugate(3,4,1,_b0,{Sigma,m},_Phi0)
                                     (3)
                         {Sigma,m} ~ iwishart(3,5,_Sigma0)
------------------------------------------------------------------------------
```

bayes: var

```
(1) Parameters are elements of the linear form xb_ogap.
(2) Parameters are elements of the linear form xb_inflation.
(3) Parameters are elements of the linear form xb_fedfunds.

Bayesian vector autoregression                  MCMC iterations  =     12,500
Gibbs sampling                                  Burn-in          =      2,500
                                                MCMC sample size =     10,000
Sample: 1956q3 thru 2005q1                      Number of obs    =        195
                                                Acceptance rate  =          1
                                                Efficiency:  min =      .9655
                                                             avg =      .9969
Log marginal-likelihood = -684.68843                         max =          1
```

bayes: var

	Mean	Std. dev.	MCSE	Median	Equal-tailed [95% cred. interval]	
ogap						
ogap						
L1.	1.031375	.044357	.000437	1.031234	.9444055	1.117762
L2.	-.0528438	.0405805	.000406	-.0527833	-.1326714	.0255082
L3.	-.047037	.0284115	.000289	-.0468732	-.1034755	.0074646
L4.	-.023866	.0217173	.000217	-.0238532	-.0664774	.0189295
inflation						
L1.	-.0671996	.062533	.000625	-.067108	-.190173	.0556707
L2.	.0055621	.0621467	.000621	.0063042	-.1185189	.1254208
L3.	.0105986	.0423772	.000429	.010359	-.0718529	.0939237
L4.	.0131083	.031908	.000311	.012586	-.049442	.0755057
fedfunds						
L1.	-.0069443	.0412664	.000413	-.0069027	-.088611	.0734544
L2.	-.0520572	.0364992	.000365	-.0522726	-.123415	.0196322
L3.	.0107723	.0261845	.000262	.0110532	-.0409668	.0613639
L4.	.0034257	.0202103	.000202	.0034853	-.0362366	.0430734
_cons	.3722142	.1206583	.001207	.373689	.1360688	.6106045

bayes: var

inflation						
ogap						
L1.	.0640502	.0300904	.000301	.0644196	.0046756	.1236306
L2.	.0079184	.027372	.000274	.0076928	-.0456988	.0619579
L3.	-.0019915	.0190604	.000191	-.0020367	-.0392003	.0355392
L4.	-.0095053	.0143533	.000139	-.009651	-.0373918	.0181917
inflation						
L1.	1.103588	.0410647	.000411	1.103125	1.023206	1.184515
L2.	-.0600424	.0410225	.000417	-.0596783	-.1406554	.0186434
L3.	-.0360733	.0285733	.000286	-.0358404	-.0917614	.019932
L4.	-.0396264	.0217553	.000218	-.0395052	-.0828371	.0034681
fedfunds						
L1.	.0763831	.0278869	.000272	.0766297	.021265	.1310788
L2.	-.0359797	.0245725	.000246	-.0358554	-.0835116	.0131129
L3.	-.0155223	.0176495	.000176	-.0154691	-.0500153	.0189255
L4.	-.0197184	.0135665	.000136	-.0197725	-.045819	.0075238
_cons	.1370817	.0806153	.000806	.137038	-.0208748	.293633

bayes: var

fedfunds						
ogap						
L1.	.1895978	.0501538	.000502	.1900072	.0902037	.2874039
L2.	-.0555335	.0458472	.000458	-.0556549	-.1457215	.0349435
L3.	-.049423	.0319737	.00032	-.0495487	-.1115743	.0134532
L4.	-.0328223	.0242815	.000243	-.0325965	-.0803066	.0147402
inflation						
L1.	.0586924	.0701206	.000701	.0586515	-.0803843	.1955202
L2.	.0556173	.0695466	.000695	.0557421	-.081407	.1931673
L3.	-.0024193	.0482795	.000483	-.002712	-.0948545	.0920376
L4.	-.0168845	.0363626	.000353	-.0167124	-.0889465	.0540924
fedfunds						
L1.	.9616559	.0467499	.000467	.9612515	.870149	1.05403
L2.	-.0726135	.0412947	.000413	-.0723247	-.1544352	.0082702
L3.	.0160287	.0294622	.000295	.0160759	-.0404996	.0744029
L4.	.001778	.0227218	.00023	.0022371	-.0427308	.0464566
_cons	.1880803	.1345058	.001345	.1873232	-.0767662	.4533603

bayes: var

Sigma_1_1	.6426288	.0653287	.000653	.6383182	.5265041	.7844193
Sigma_2_1	.0255611	.0309588	.000314	.025407	-.0348323	.087608
Sigma_3_1	.2353097	.0539973	.00054	.2332346	.1358691	.3470781
Sigma_2_2	.2887829	.0291186	.000287	.2870285	.2366544	.3507061
Sigma_3_2	.1433862	.0360754	.000361	.142147	.0768442	.2192838
Sigma_3_3	.8167953	.0827291	.000827	.8111552	.6704141	.9966348

(output completes)

VAR postestimation

- VAR coefficients are not usually interesting by themselves
- We are usually interested instead in
    - VAR stability
    - forecasting
    - impulse response analysis
- ...which are functions of the VAR coefficients

Bayesian VAR postestimation

- Specialized postestimation for Bayesian VARs:
  - `bayesvarstable`
  - `bayesfcast`
  - `bayesirf`
- General Bayesian postestimation features
  - `bayesstats grubin`
  - `bayesstats ppvalues`
  - `bayesstats summary`
  - `bayesstats ess`
  - `bayespredict`
  - `bayesgraph`

VAR stability

- A VAR is said to be stable if the eigenvalues of its companion matrix are all strictly less than 1 in absolute value

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 & \dots & \mathbf{A}_{p-1} & \mathbf{A}_p \\ \mathbf{I} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & \dots & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{I} & \mathbf{0} \end{bmatrix} \qquad (kp \times kp)$$

- We can compute the whole posterior distribution of these eigenvalues
- Stata: `bayesvarstable`

bayesvarstable

```
. bayes, saving(bvar.dta, replace)
note: file bvar.dta saved.

. bayesvarstable

Eigenvalue stability condition                 Companion matrix size =    12
                                               MCMC sample size      = 10000
```

Eigenvalue modulus	Mean	Std. dev.	MCSE	Median	Equal-tailed [95% cred.	interval]
1	.9475916	.0193222	.000193	.9485289	.9070583	.982624
2	.9425437	.0248523	.000249	.9457567	.8804373	.9803136
3	.8237587	.0711348	.000711	.8336736	.6805061	.9332017
4	.5891391	.0934168	.000934	.5782689	.4227001	.7720391
5	.4828312	.0890523	.000891	.4834793	.3297175	.6502389
6	.3671857	.0426969	.000427	.3645048	.2903083	.4624801
7	.3505756	.0366104	.000366	.3509496	.2778521	.422465
8	.3172101	.0378505	.000379	.3190526	.2387078	.3879864
9	.3026827	.0387949	.000388	.304546	.2214313	.3724483
10	.2672503	.0472622	.000473	.2719659	.1628997	.3458558
11	.236392	.0548257	.000548	.2420806	.1161736	.3293138
12	.1906105	.0789222	.000789	.2054165	.0169675	.3097373

```
Pr(eigenvalues lie inside the unit circle) = 0.9979
```

VAR forecasting

- We compute a forecast $(\widehat{\mathbf{y}}^s_{T+1}, \ldots, \widehat{\mathbf{y}}^s_{T+h})$ from each of the s MCMC samples, $\theta^s = (\beta^s, \mathbf{\Sigma}^s)$
- . . . to arrive at the whole posterior distribution of forecasts
- Stata: `bayesfcast compute` and `bayesfcast graph`

bayesfcast

```
. bayesfcast compute bf_, step(12)
. bayesfcast graph bf_inflation, observed
```

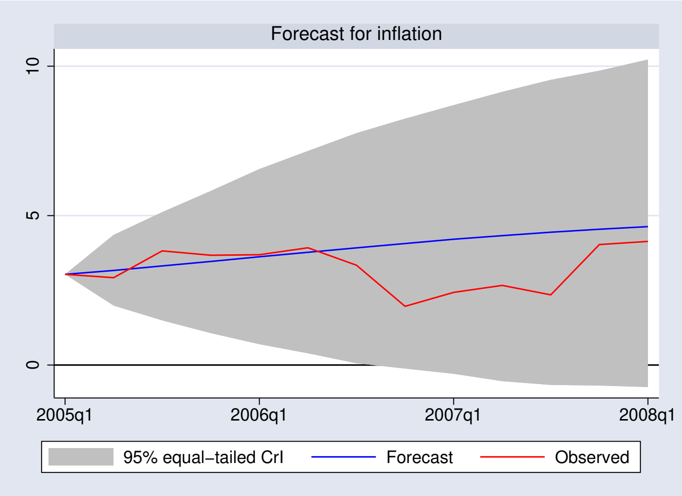

Impulse response functions

- Impulse response functions trace out how a shock to one equation affects model variables
- `bayesirf` is a suite of commands for creating, managing, and analyzing impulse response functions (patterned after `irf`)
  - `bayesirf set`
  - `bayesirf create`
  - `bayesirf table`
  - `bayesirf graph`
  - ...among others
- `bayesirf create` builds a collection of IRF results, including
  - simple IRF (`irf`)
  - orthogonalized IRF (`oirf`)
  - cumulative IRF (`cirf`, `coirf`)
  - forecast error variance decomposition (`fevd`)
  - dynamic multiplier (`dm`) in the presence of exogenous variables

bayesirf

```
. bayesirf set bvarirf.irf, replace
(file bvarirf.irf created)
(file bvarirf.irf now active)
. bayesirf create model1, step(20)
(file bvarirf.irf updated)
. bayesirf graph oirf, yline(0, lcolor(black))
```

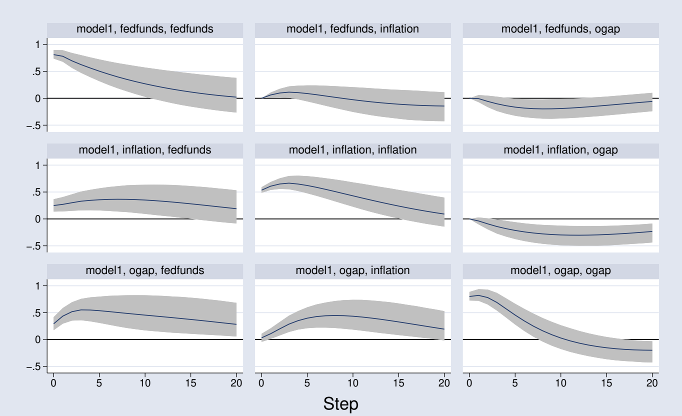

bayes: var options

- Controlling the prior: `minnconjprior()`
  - `selftight(#)`
  - `crosstight(#)`
  - `lagdecay(#)`
  - `mean(#)`

`bayes: dsge`

Dynamic Stochastic General Equilibrium models

- Vector autoregression models have a minimum of structure
  - choose variables and lag length, and perhaps order
- Dynamic stochastic general equilibrium models have lots of structure
  - n variables in n equations
  - Equations can feature lags and *leads*
  - Some components are latent (unobserved)
  - Equations are motivated by economic theory
- DSGE models are solved into state-space form and estimated based on the likelihood of the state-space solution

A simple DSGE model

- A model with 3 control variables, driven by 2 state variables
- Equations:

$$\begin{aligned} x_t &= E_t x_{t+1} - \sigma(r_t - p_{t+1} - z_t) \\ p_t &= \beta E_t p_{t+1} + \kappa x_t \\ r_t &= \frac{1}{\psi} p_t + w_t \\ z_{t+1} &= \rho_z z_t + \epsilon_{t+1} \\ w_{t+1} &= \rho_w w_t + e_{t+1} \end{aligned}$$

- x_t is the output gap, p_t is inflation, r_t is the interest rate
- z_t and w_t are driving state variables
- e_t and ϵ_t are shocks

Priors for DSGE model parameters

- VARs tend to have many parameters
- DSGEs tend to be tightly parameterized, with parameters that have immediate theoretical interpretation
  - Benefit: smaller set of parameters to estimate
  - Cost: you lean heavily on theory and model specification
- Many DSGE model parameters have natural bounds or theoretical considerations that provide useful priors
  - With AR(1) state variables, autoregressive parameters must lie in (-1,1) for stability
  - Many parameters represent shares or rates that must lie in (0,1)
  - Beta distributions lie in (0,1) and give extra weight to specific parts of that interval, making them a popular choice

bayes: dsge

```
. webuse usmacro2

. bayes, prior({beta},  beta(95, 5))   prior({kappa}, beta(30,70))     ///
>       prior({sigma}, beta(10,90))   prior({psi},   beta(67,33))     ///
>       prior({rhow},  beta(10, 10))  prior({rhoz},  beta(35,15))     ///
>       rseed(17) dots burnin(5000) mcmcsize(30000):                  ///
>       dsge (x   = F.x - {sigma}*(r - F.p - z)    , unobserved ) ///
>            (p   = {beta}*F.p + {kappa}*x                     ) ///
>            (r   = (1/{psi})*p + w                            ) ///
>            (F.z = {rhoz}*z                  , state          ) ///
>            (F.w = {rhow}*w                  , state          )
note: initial parameter vector set to means of priors.

Burn-in 5000 aaaaaaaaa1000aaaaaaaaa2000aaaa.....3000.........4000.........5000
> done
Simulation 30000 .........1000.........2000.........3000.........4000.........5
> 000.........6000.........7000.........8000.........9000.........10000........
> .11000.........12000.........13000.........14000.........15000.........16000.
> ........17000.........18000.........19000.........20000.........21000........
> .22000.........23000.........24000.........25000.........26000.........27000.
> ........28000.........29000.........30000 done
```

bayes: dsge

```
Model summary
------------------------------------------------------------------------------
Likelihood:
  p r ~ dsgell({sigma},{beta},{kappa},{psi},{rhoz},{rhow},{sd(e.z)},{sd(e.w)})

Priors:
            {sigma} ~ beta(10,90)
             {beta} ~ beta(95,5)
            {kappa} ~ beta(30,70)
              {psi} ~ beta(67,33)
             {rhoz} ~ beta(35,15)
             {rhow} ~ beta(10,10)
  {sd(e.z) sd(e.w)} ~ igamma(.01,.01)
------------------------------------------------------------------------------
```

bayes: dsge

```
Bayesian linear DSGE model                         MCMC iterations  =     35,000
Random-walk Metropolis-Hastings sampling           Burn-in          =      5,000
                                                   MCMC sample size =     30,000
Sample: 1955q1 thru 2015q4                         Number of obs    =        244
                                                   Acceptance rate  =      .1741
                                                   Efficiency:  min =    .005331
                                                                avg =     .01032
Log marginal-likelihood = -787.73905                            max =     .01974
```

	Mean	Std. dev.	MCSE	Median	Equal-tailed [95% cred.	interval]
sigma	.1443227	.0292876	.001318	.1432416	.088498	.2043906
beta	.9547238	.0203592	.001276	.9576523	.9077723	.9848212
kappa	.3419745	.0457376	.003318	.3415295	.2517864	.4357346
psi	.6527897	.043041	.003403	.6529768	.5686343	.7351054
rhoz	.9078086	.0157278	.00091	.9080455	.8749843	.9369648
rhow	.7546737	.0269813	.001109	.7541327	.7017534	.8085894
sd(e.z)	.6048148	.0950875	.005623	.5951787	.4475909	.8237128
sd(e.w)	1.955303	.1265823	.008904	1.948905	1.734296	2.230285

Bayesian DSGE postestimation

- Specialized postestimation for Bayesian DSGEs:
  - `bayesirf`
- General Bayesian postestimation features
  - `bayesstats grubin`
  - `bayesstats ppvalues`
  - `bayesstats summary`
  - `bayesstats ess`
  - `bayesgraph`

Posterior parameter diagnostic plots

```
. bayesstats summary (1/{psi})

Posterior summary statistics                      MCMC sample size =     30,000

      expr1 : 1/{psi}
```

	Mean	Std. dev.	MCSE	Median	Equal-tailed [95% cred.	interval]
expr1	1.538681	.1035836	.008064	1.531448	1.360349	1.7586

```
. bayesgraph diagnostics (1/{psi})
```

Posterior parameter diagnostic plots

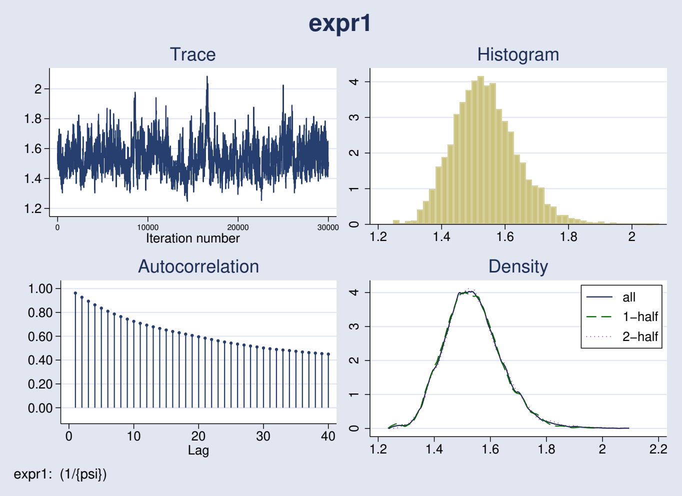

Posterior parameter distribution plots

```
. bayesgraph kdensity {kappa}, lcolor(red) ///
>         addplot(function Prior=betaden(30,70,x), ///
>         legend(on label(1 "Posterior")) lcolor(blue)) name(kappa) nodraw

.
. bayesgraph kdensity {sigma}, lcolor(red) ///
>         addplot(function Prior=betaden(10,90,x), ///
>         legend(on label(1 "Posterior")) lcolor(blue)) name(sigma) nodraw

.
. bayesgraph kdensity {beta}, lcolor(red) ///
>         addplot(function Prior=betaden(95,5,x), ///
>         legend(on label(1 "Posterior")) lcolor(blue)) name(beta) nodraw

.
. bayesgraph kdensity {psi}, lcolor(red) ///
>         addplot(function Prior=betaden(67,33,x), ///
>         legend(on label(1 "Posterior")) lcolor(blue)) name(psi) nodraw

.
. bayesgraph kdensity {rhoz}, lcolor(red) ///
>         addplot(function Prior=betaden(35,15,x), ///
>         legend(on label(1 "Posterior")) lcolor(blue)) name(rhoz) nodraw

.
. bayesgraph kdensity {rhow}, lcolor(red) ///
>         addplot(function Prior=betaden(10,10,x), ///
>         legend(on label(1 "Posterior")) lcolor(blue)) name(rhow) nodraw

.
. graph combine kappa sigma beta psi rhoz rhow
```

Posterior parameter distribution plots

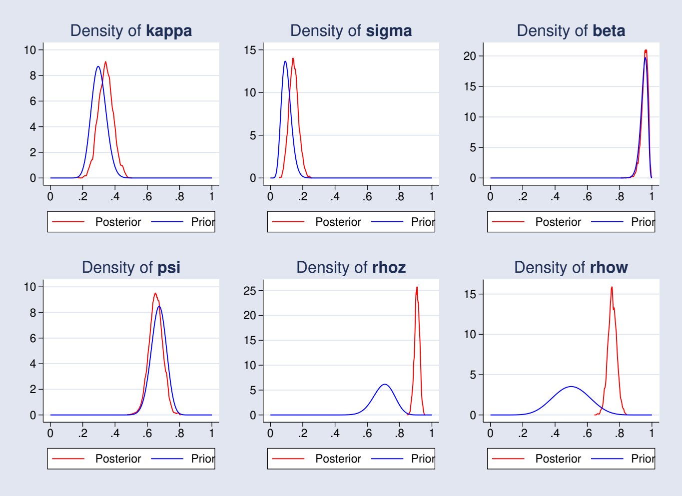

Impulse response functions

```
. bayesirf set bdsgeirf.irf, replace
(file bdsgeirf.irf created)
(file bdsgeirf.irf now active)

. bayesirf create model1
(file bdsgeirf.irf updated)

. bayesirf graph irf, impulse(w) response(p x r w)   ///
>         byopts(yrescale) yline(0, lcolor(black))
```

Impulse response functions

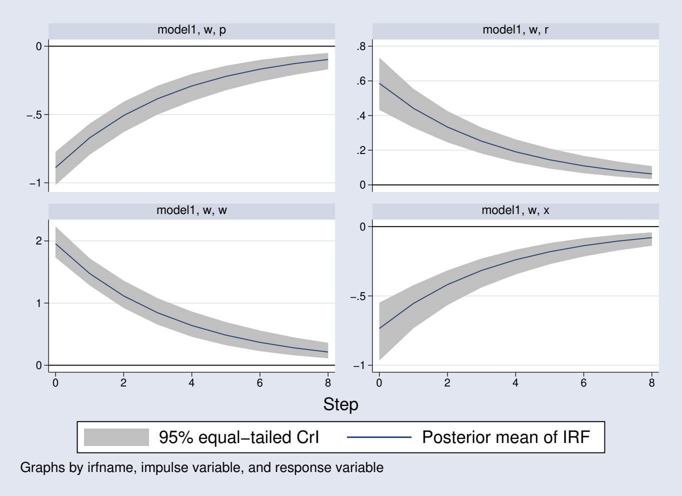

`bayes: dsgenl`

Nonlinear DSGE models

- DSGE models in which the equations are nonlinear in variables
- Many theories come in nonlinear form
- The solution must be approximated
- Approximation technique: first-order perturbation
- Solution: linearized state-space form
- Estimation: maximum likelihood or Bayesian

A nonlinear DSGE model

- A model with 3 control variables, driven by 2 state variables
- Equations:

$$\frac{1}{c_t} = \beta E_t \left[\left(\frac{1}{c_{t+1}} \right) (1 + r_{t+1} - \delta) \right]$$

$$r_t = \alpha \frac{y_t}{k_t}$$

$$y_t = z_t k_t^{\alpha}$$

$$k_{t+1} = y_t - c_t + (1 - \delta) k_t$$

$$\ln(z_{t+1}) = \rho \ln z_t + e_{t+1}$$

- Control variables: y_t, c_t, r_t
- State variables: k_t, z_t
- Shock: e_t

Priors for DSGE models

- Similar considerations as for linear DSGE models
- This model has several parameters with interpretations:
  - α is the capital share of income, roughly 0.3 and in $(0, 1)$
  - δ is the depreciation rate, roughly 0.05 and in $(0, 1)$
  - β is the discount factor, roughly 0.96 and in $(0, 1)$
  - ρ is an autoregressive parameter, likely positive and in $(0, 1)$

bayes: dsgenl

```
. webuse usmacro2

. bayes, prior({alpha}, beta(30,70))  prior({beta}, beta(95,5))    ///
>        prior({delta}, beta(25,975))  prior({rho}, beta(5, 3))    ///
>        rseed(17) burnin(5000) dots :                             ///
>        dsgenl (1/c       = {beta}*(1/f.c)*(1+f.r-{delta}))           ///
>                (r        = {alpha}*y/k)                               ///
>                (y        = z*k^({alpha}))                             ///
>                (f.k      = y - c + (1-{delta})*k)                     ///
>                (ln(f.z) = {rho}*ln(z))                               ///
>                , exostate(z) endostate(k) observed(y) unobserved(c r)
note: initial parameter vector set to means of priors.

Burn-in 5000 aaaaaaaaa1000aaaaaaaaa2000aaaaaaaaa3000aaaaaa...4000.........5000
> done
Simulation 10000 .........1000.........2000.........3000.........4000.........5
> 000.........6000.........7000.........8000.........9000.........10000 done
```

bayes: dsgenl

```
Model summary
------------------------------------------------------------------------------
Likelihood:
  y ~ dsgell({beta},{delta},{alpha},{rho},{sd(e.z)})

Priors:
     {beta} ~ beta(95,5)
    {delta} ~ beta(25,975)
    {alpha} ~ beta(30,70)
      {rho} ~ beta(5,3)
  {sd(e.z)} ~ igamma(.01,.01)
------------------------------------------------------------------------------
```

bayes: dsgenl

```
Bayesian first-order DSGE model                     MCMC iterations  =     15,000
Random-walk Metropolis-Hastings sampling            Burn-in          =      5,000
                                                    MCMC sample size =     10,000
Sample: 1955q1 thru 2015q4                          Number of obs    =        244
                                                    Acceptance rate  =      .2506
                                                    Efficiency:  min =     .04563
                                                                 avg =     .04999
Log marginal-likelihood = -649.82949                             max =     .05552
```

	Mean	Std. dev.	MCSE	Median	Equal-tailed [95% cred.	interval]
beta	.9554009	.0194803	.00086	.9582636	.9104263	.9847218
delta	.0250477	.0048846	.000207	.0247132	.0163193	.0357851
alpha	.2962864	.0442748	.002073	.2958201	.2122214	.3837115
rho	.3064437	.0584439	.002654	.3047101	.1939408	.422685
sd(e.z)	3.359195	.152429	.006891	3.360512	3.068458	3.658345

bayes: dsgenl prior–posterior plots

```
. bayesgraph kdensity {alpha}, lcolor(red) ///
>         addplot(function Prior=betaden(30,70,x), ///
>         legend(on label(1 "Posterior")) lcolor(blue)) name(alpha) nodraw

.
. bayesgraph kdensity {beta}, lcolor(red) ///
>         addplot(function Prior=betaden(95, 5,x), ///
>         legend(on label(1 "Posterior")) lcolor(blue)) name(beta) nodraw

.
. bayesgraph kdensity {delta}, lcolor(red) ///
>         addplot(function Prior=betaden(25,975,x), ///
>         legend(on label(1 "Posterior")) lcolor(blue)) name(delta) nodraw

.
. bayesgraph kdensity {rho}, lcolor(red) ///
>         addplot(function Prior=betaden(5,3,x), ///
>         legend(on label(1 "Posterior")) lcolor(blue)) name(rho) nodraw

.
. graph combine alpha beta delta rho
```

bayes: dsgenl prior–posterior plots

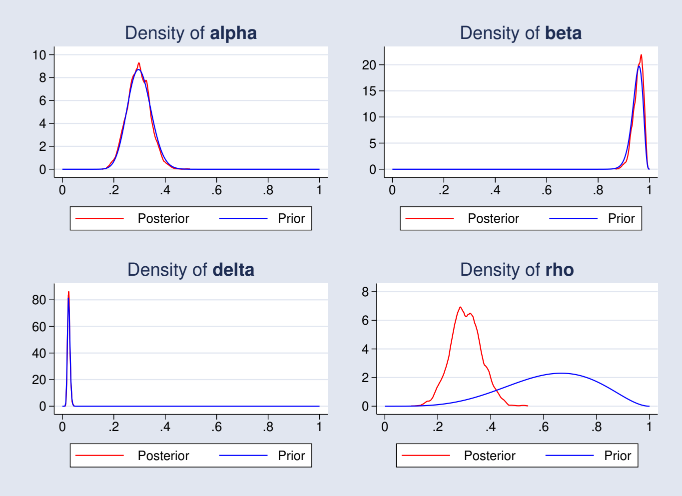

bayes: dsgenl IRFs

```
. bayesirf set stochmodel, replace
(file stochmodel.irf created)
(file stochmodel.irf now active)

. bayesirf create model1
(file stochmodel.irf updated)

. bayesirf graph irf, impulse(z) byopts(yrescale)
```

bayes: dsgenl IRFs

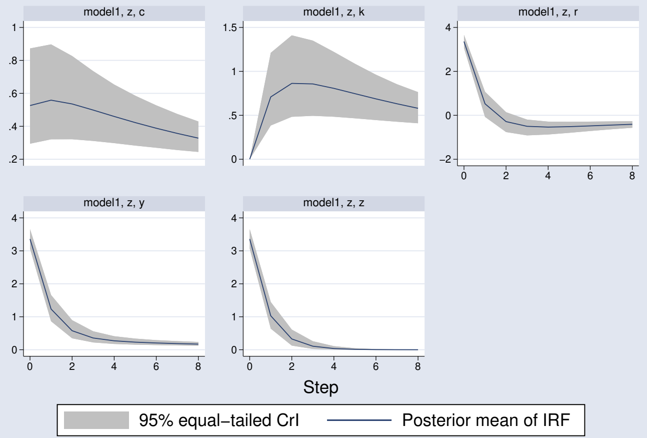

Graphs by irfname, impulse variable, and response variable `bayes: xt`

Bayesian panel–data commands

- The `bayes:` prefix also works with many of Stata's panel–data commands
  - `bayes: xtreg`
  - `bayes: xtlogit`
  - `bayes: xtprobit`
  - `bayes: xtologit`
  - `bayes: xtoprobit`
  - `bayes: xtmlogit`
  - `bayes: xtpoisson`
  - `bayes: xtnbreg`

bayes: xtreg

```
. webuse nlswork6
(Subsample of 1986 National Longitudinal Survey of Young Women)

. xtset

Panel variable: id (unbalanced)
 Time variable: year, 68 to 88, but with gaps
         Delta: 1 unit

. describe id year ln_wage grade ttl_exp not_smsa

Variable      Storage   Display    Value
    name         type    format    label      Variable label
-------------------------------------------------------------------------------

id              int     %9.0g                 ID
year            byte    %8.0g                 Interview year
ln_wage         float   %9.0g                 ln(wage/GNP deflator)
grade           byte    %8.0g                 Current grade completed
ttl_exp         float   %9.0g                 Total work experience
not_smsa        byte    %8.0g                 1 if not SMSA
```

bayes: xtreg

```
. bayes, rseed(17): xtreg ln_wage grade ttl_exp i.not_smsa
note: Gibbs sampling is used for regression coefficients and variance
      components.

Burn-in 2500 aaaaaaaaa1000aaaaaaaaa2000aaaaa done
Simulation 10000 .........1000.........2000.........3000.........4000.........5
> 000.........6000.........7000.........8000.........9000.........10000 done
Model summary
------------------------------------------------------------------------------
Likelihood:
  ln_wage ~ normal(xb_ln_wage,{sigma2})

Priors:
  {ln_wage:grade ttl_exp 1.not_smsa _cons} ~ normal(0,10000)                (1)
                                  {U[id]} ~ normal(0,{var_U})               (1)
                                 {sigma2} ~ igamma(0.01,0.01)

Hyperprior:
  {var_U} ~ igamma(0.01,0.01)
------------------------------------------------------------------------------
(1) Parameters are elements of the linear form xb_ln_wage.
```

bayes: xtreg

```
Bayesian RE normal regression                    MCMC iterations  =     12,500
Metropolis-Hastings and Gibbs sampling           Burn-in          =      2,500
                                                 MCMC sample size =     10,000
Group variable: id                               Number of groups =        831
                                                 Obs per group:
                                                              min =          1
                                                              avg =        1.4
                                                              max =          5
                                                 Number of obs    =      1,174
                                                 Acceptance rate  =      .7975
                                                 Efficiency:  min =     .02034
                                                              avg =     .05583
Log marginal-likelihood                                       max =      .1233
```

bayes: xtreg

	Mean	Std. dev.	MCSE	Median	Equal-tailed [95% cred.	interval]
ln_wage						
grade	.0704289	.0052453	.000246	.0702965	.060247	.0810986
ttl_exp	.0320736	.0022917	.000065	.0320706	.0276516	.0366477
1.not_smsa	-.1453248	.026093	.000979	-.1452854	-.1964133	-.0945867
_cons	.5748128	.0653271	.002939	.5766947	.4437356	.7002126
var_U	.0807316	.0076072	.000533	.0804105	.0668763	.0960505
sigma2	.0672671	.005118	.00032	.0670958	.0577482	.0777501

Note: Default priors are used for model parameters.

bayes: xtreg postestimation

```
. bayesgraph histogram {U[1/9]}, byparm normal
```

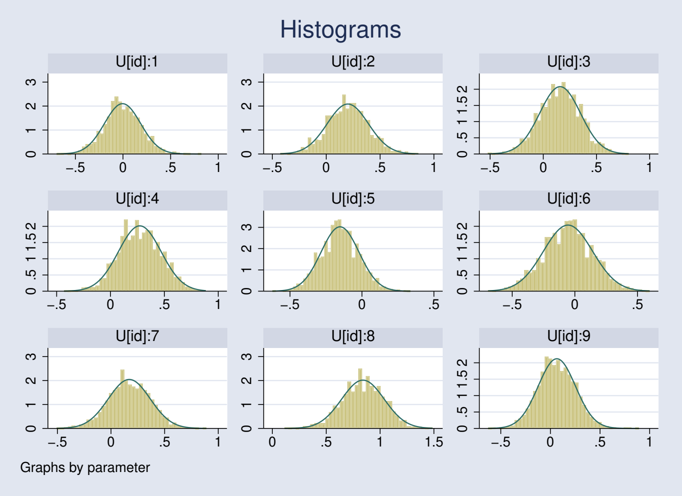

Recap

- Stata 17 introduces Bayesian estimation of a variety of time–series and panel–data econometric models
  - Multivariate time-series:
    - `bayes: var`
    - `bayes: dsge`
    - `bayes: dsgenl`
    - `bayesirf`
    - `bayesfcast`
  - Panel data:
    - `bayes:  xtreg`
    - `bayes:  xtlogit`
    - `bayes:  xtprobit`
    - `bayes:  xtologit`
    - `bayes:  xtoprobit`
    - `bayes:  xtmlogit`
    - `bayes:  xtpoisson`
    - `bayes:  xtnbreg`

Thank you!