Title

stata.com

Intro 9b — Bayesian estimation of stochastic growth model

Description   Remarks and examples   Reference   Also see

Description

This introduction estimates and interprets the parameters of a simple stochastic growth model using Bayesian methods.

Remarks and examples

stata.com

Remarks are presented under the following headings:

The model
Parameter estimation
Posterior diagnostics and plots
Impulse responses

The model

The model contains equations that jointly determine output Y_t, the interest rate R_t, consumption C_t, capital K_t, and productivity Z_t. The model contains four parameters: α, β, δ, and ρ. This model is a variant on the model used in Schmitt-Grohé and Uribe (2004):

$$
\begin{aligned}
1 &= \beta E_t \left\{ \left(\frac{C_{t+1}}{C_t} \right)^{-1} (1 + R_{t+1} - \delta) \right\} && (1) \\
Y_t &= Z_t K_t^{\alpha} && (2) \\
R_t &= \alpha Z_t K_t^{\alpha - 1} && (3) \\
K_{t+1} &= Y_t - C_t + (1 - \delta) K_t && (4) \\
\ln(Z_{t+1}) &= \rho \ln(Z_t) + e_{t+1} && (5)
\end{aligned}
$$

Equation (1) defines a relationship between consumption growth C_{t+1}/C_t and the real interest rate R_{t+1}. Equation (2) is a production function for output Y_t as a function of productivity Z_t and capital K_t. Equation (3) is a model for the interest rate. Equation (4) is the equation for capital accumulation; capital in the next period is equal to underappreciated capital this period $(1-\delta)K_t$ plus unconsumed output $Y_t - C_t$. Equation (5) is a law of motion for productivity Z_t. The parameter β is a discount factor in the consumption equation, the parameter α is a production parameter in the output equation, the parameter δ is a depreciation parameter in the capital equation, and the parameter ρ is a persistence parameter in the productivity equation.

The state variables are the current-period capital stock and the level of productivity, (K_t, Z_t). The control variables are consumption, the interest rate, and output, (C_t, R_t, Y_t). We estimate the parameters of the linearized version of the model using Bayesian methods.

Parameter estimation

We specify priors for the model parameters. The discount rate β must lie between 0 and 1, with common values in the range (0.95, 0.99). The capital share α must lie between 0 and 1 and is usually estimated to be about one-third. The depreciation rate δ must lie between 0 and 1, and a common value for it is 0.025 or 0.10. The autocorrelation parameter ρ must be less than 1 in absolute value and is usually thought to be positive and close to 1. Our prior choices for all parameters are driven by these theoretical considerations. As all four parameters are plausibly restricted to the unit interval, a beta distribution is chosen for all four priors.

The parameters of the beta distribution were chosen to put the weight of prior mass on theoretically appropriate values. For the discount factor `{beta}`, this is the range (0.90, 0.99). For the depreciation parameter `{delta}`, this is the range (0.03, 0.05). For the capital share `{alpha}`, this is the range (0.3, 0.4). For the autoregressive parameter, this is the range (0.6, 0.99). The prior means for each parameter are as follows: `{beta}` prior mean is 0.95; `{delta}` prior mean is 0.044; `{alpha}` prior mean is 0.38; and `{rhoz}` prior mean is 0.8.

```
. use https://www.stata-press.com/data/r18/usmacro2
(Federal Reserve Economic Data - St. Louis Fed, 2017-01-15)
. bayes, prior({beta},  beta( 95,   5))
>        prior({delta}, beta( 40, 860))
>        prior({alpha}, beta(360, 590))
>        prior({rhoz},  beta( 80,  20)) dots rseed(17) :
>        dsgenl (1 = {beta}*(c/F.c)*(1 + F.r - {delta}))
>               (y = z*k^({alpha}))
>               (r = {alpha}*y/k)
>               (F.k = y - c + (1-{delta})*k)
>               (ln(F.z) = {rhoz}*ln(z)),
>               exostate(z) endostate(k) observed(y) unobserved(c r)
note: initial parameter vector set to means of priors.
Burn-in 2500 aaaaaaaaa1000aaaaaaaaa2000..... done
Simulation 10000 .........1000.........2000.........3000.........4000.........
> 5000.........6000.........7000.........8000.........9000.........10000 done
Model summary
------------------------------------------------------------------------------
Likelihood:
  y ~ dsgell({beta},{delta},{alpha},{rhoz},{sd(e.z)})
Priors:
     {beta} ~ beta(95,5)
    {delta} ~ beta(40,860)
    {alpha} ~ beta(360,590)
     {rhoz} ~ beta(80,20)
  {sd(e.z)} ~ igamma(.01,.01)
------------------------------------------------------------------------------
```

```
Bayesian first-order DSGE model                  MCMC iterations  =     12,500
Random-walk Metropolis-Hastings sampling         Burn-in          =      2,500
                                                 MCMC sample size =     10,000
Sample: 1955q1 thru 2015q4                       Number of obs    =        244
                                                 Acceptance rate  =      .2479
                                                 Efficiency:  min =      .0228
                                                              avg =        .05
Log marginal-likelihood = -670.33605                          max =     .05948

                                                            Equal-tailed
              Mean   Std. dev.     MCSE     Median  [95% cred. interval]

     beta  .9696788   .0138534   .000918  .9714437   .9379939   .9904246
    delta  .0412205   .0063456   .000272  .0408097    .029939   .0545771
    alpha  .3790517   .0154948   .000645  .3792245   .3486508   .4096519
     rhoz  .6178456   .0442668   .001815   .618141   .5340957   .7060575
  sd(e.z)   3.54923   .1751019   .007431  3.544049   3.238064   3.931496
```

The model summary reports the prior and likelihood specifications, including the default inverse-gamma prior for the standard deviation of the shock.

The output header reports the burn-in length and MCMC sample size, as well as information about the efficiency of the Metropolis–Hastings sampler. The overall acceptance rate is 0.25, with sampling efficiencies between 0.023 and 0.059.

The posterior mean for `{beta}` is 0.97, close to the prior mean of 0.95. The posterior mean for `{delta}` is 0.041, close to its prior mean of 0.044. The posterior mean for `{alpha}` is 0.38, identical to its prior mean. The posterior mean for `{rhoz}` is 0.62, substantially different from its prior mean of 0.80. Overall, many of the parameters show little updating, indicating that the likelihood is uninformative along several dimensions of the model's parameter space. The posterior results for `{beta}`, `{delta}`, and `{alpha}` are mainly driven by the prior.

Posterior diagnostics and plots

We begin by investigating effective sample sizes for each parameter.

```
. bayesstats ess

Efficiency summaries    MCMC sample size =    10,000
                        Efficiency:  min =     .0228
                                     avg =       .05
                                     max =    .05948

                 ESS   Corr. time    Efficiency

     beta     227.97        43.87        0.0228
    delta     545.58        18.33        0.0546
    alpha     576.58        17.34        0.0577
     rhoz     594.79        16.81        0.0595
  sd(e.z)     555.25        18.01        0.0555
```

The effective sample size for the discount factor `{beta}` is somewhat low relative to the other parameters, which indicates that blocking may improve sampling efficiency.

Because {rhoz} was the only internal parameter to receive substantial updating, we look at its full set of posterior diagnostic plots.

```
. bayesgraph diagnostics {rhoz}
```

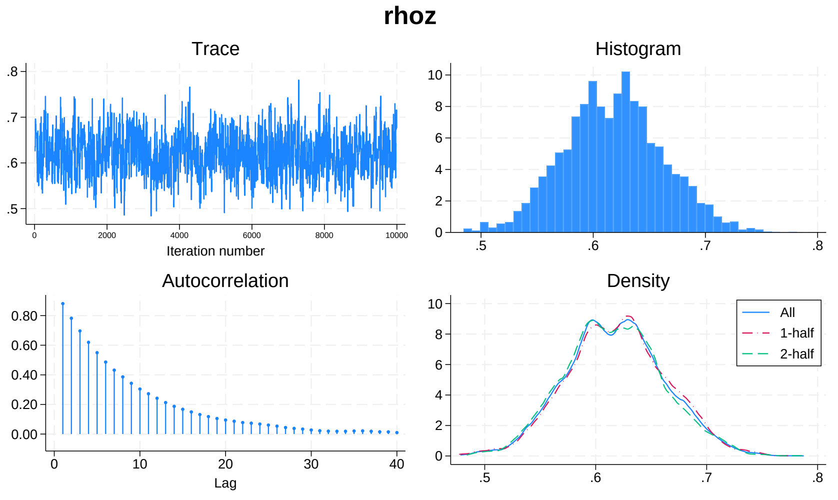

Autocorrelations tail off at a moderate pace, the trace plot shows reasonable mixing, and the density plot shows that the first- and second-half densities do not substantially differ from the full-sample density.

Next, we generate prior–posterior plots for two parameters, the capital share and the autoregressive parameter.

```
. bayesgraph kdensity {alpha},
> addplot(function Prior = betaden(360, 590, x),
> legend(on label(1 "Posterior")))
. bayesgraph kdensity {rhoz},
> addplot(function Prior = betaden(80, 20, x),
> legend(on label(1 "Posterior")))
```

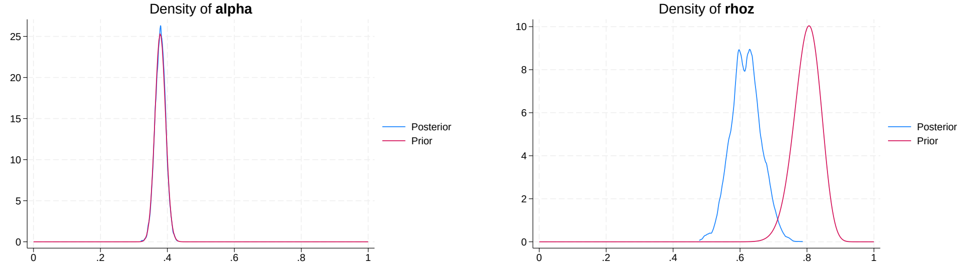

The posterior density of {alpha} does not differ from its prior density. This situation indicates a flat likelihood along the {alpha} dimension. By contrast, the posterior density of {rhoz} does differ from its prior density. The posterior mean has fallen to 0.6 from 0.8.

Prior–posterior plots for the discount rate {beta} and the depreciation parameter {delta} are

```
. bayesgraph kdensity {beta},
> addplot(function Prior = betaden(95, 5, x),
> legend(on label(1 "Posterior")))
. bayesgraph kdensity {delta},
> addplot(function Prior = betaden(40, 860, x),
> legend(on label(1 "Posterior")))
```

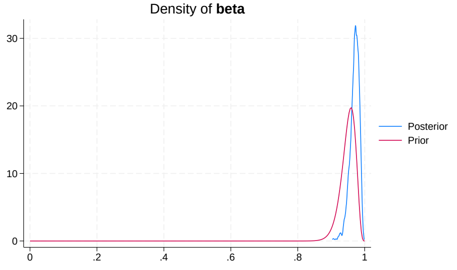

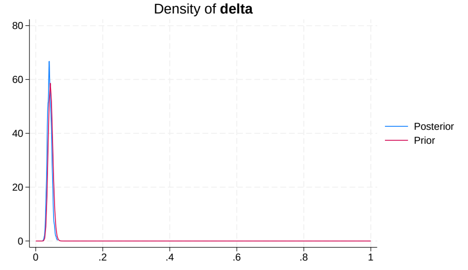

The posterior distribution of {beta} lies close to its prior. The posterior distribution of {delta} lies on top of its prior, indicating that the data provide little information along this dimension of the model.

Impulse responses

Next, we explore the response of model variables to a shock to the state variable z. We begin by saving off our MCMC results in a dataset.

```
. bayes, saving(bayes_dsgenl_sim, replace)
```

Next, we set up the impulse–response function file and impulse–response functions themselves. We are using the [BAYES] **bayesirf** command. We specify step(20) to plot the first 20 periods after the shock. Because the unit of time in this model is one quarter, 20 periods correspond to 5 years.

```
. bayesirf set rbcmodel
. bayesirf create model1, step(20)
. bayesirf graph irf, impulse(z) byopts(yrescale)
```

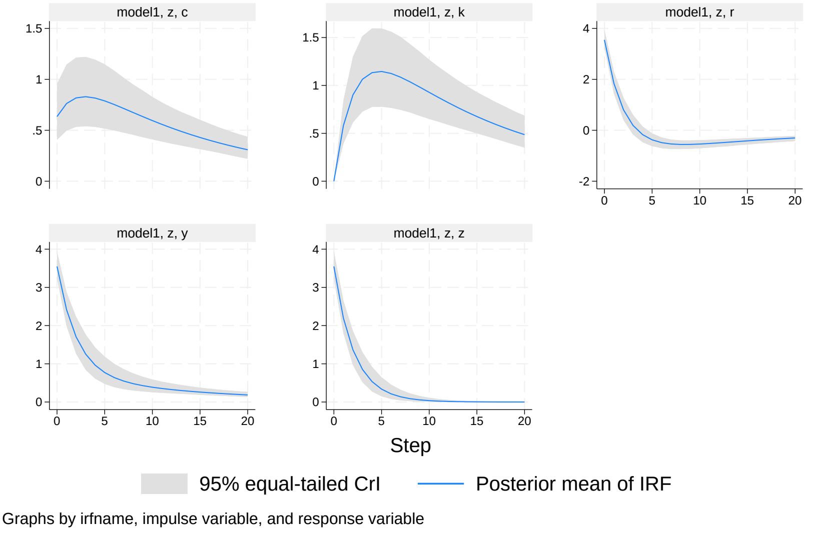

Each panel displays the response of one model variable to the impulse. Each step is one quarter, so that four steps are one year after the shock. In the top-left panel, consumption `c` rises and follows a mostly flat trajectory for the first eight periods after the shock and then falls to return to its steady state. In the top-middle panel, the capital stock `k` does not move in the first period but rises afterward in a hump-shaped pattern. In the top-right panel, the interest rate `r` rises on impact, remains elevated for the first four periods, and then dips below its steady-state value in the fifth period; it returns to its steady state from below. In the bottom-left panel, output `y` rises on impact and then declines monotonically back to its steady state. The bottom-right panel shows the evolution of the state variable `z` itself; it rises on a shock and then falls monotonically back to its steady state.

Reference

Schmitt-Grohé, S., and M. Uribe. 2004. Solving dynamic general equilibrium models using a second-order approximation to the policy function. *Journal of Economic Dynamics and Control* 28: 755–775. https://doi.org/10.1016/S0165-1889(03)00043-5.

Also see

[DSGE] **dsgenl** — Nonlinear dynamic stochastic general equilibrium models

[DSGE] **Intro 1** — Introduction to DSGEs

[DSGE] **Intro 3d** — Nonlinear New Keynesian model

[DSGE] **Intro 3f** — Stochastic growth model

[BAYES] **bayes: dsgenl** — Bayesian nonlinear dynamic stochastic general equilibrium models

[BAYES] **bayes: dsge postestimation** — Postestimation tools for bayes: dsge and bayes: dsgenl

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

For suggested citations, see the FAQ on citing Stata documentation.