

# Bayesian Inference for DSGE Models

Lawrence J. Christiano

# Outline

- Bayes' rule.
- Monte Carlo integration: a simple example.
- Markov Chain Monte Carlo (MCMC) algorithm.
- Laplace approximation

# Bayesian Inference

- Bayesian inference is about describing the mapping from prior beliefs about  $\theta$ , summarized in  $p(\theta)$ , to new posterior beliefs in the light of observing the data,  $Y^{data}$ .
- General property of probabilities:

$$p(Y^{data}, \theta) = \begin{cases} p(Y^{data}|\theta) \times p(\theta) \\ p(\theta|Y^{data}) \times p(Y^{data}) \end{cases},$$

which implies Bayes' rule:

$$p(\theta|Y^{data}) = \frac{p(Y^{data}|\theta) p(\theta)}{p(Y^{data})},$$

mapping from prior to posterior induced by  $Y^{data}$ .

# Bayesian Inference

- Report features of the posterior distribution,  $p(\theta|Y^{data})$ .
  - The value of  $\theta$  that maximizes  $p(\theta|Y^{data})$ , 'mode' of posterior distribution.
  - Compare marginal prior,  $p(\theta_i)$ , with marginal posterior of individual elements of  $\theta$ ,  $g(\theta_i|Y^{data})$  :

$$g(\theta_i|Y^{data}) = \int_{\theta_{j \neq i}} p(\theta|Y^{data}) d\theta_{j \neq i} \text{ (multiple integration!!)}$$

- Probability intervals about the mode of  $\theta$  ('Bayesian confidence intervals'), need  $g(\theta_i|Y^{data})$ .
- Marginal likelihood for assessing model 'fit':

$$p(Y^{data}) = \int_{\theta} p(Y^{data}|\theta) p(\theta) d\theta \text{ (multiple integration)}$$

# Monte Carlo Integration: Simple Example

- Much of Bayesian inference is about multiple integration.
- Numerical methods for multiple integration:
  - Quadrature integration (example: approximating the integral as the sum of the areas of triangles beneath the integrand).
  - Monte Carlo Integration: uses random number generator.
- Example of Monte Carlo Integration:
  - suppose you want to evaluate

$$\int_a^b f(x) dx, \quad -\infty \leq a < b \leq \infty.$$

- select a density function,  $g(x)$  for  $x \in [a, b]$  and note:

$$\int_a^b f(x) dx = \int_a^b \frac{f(x)}{g(x)} g(x) dx = E \frac{f(x)}{g(x)},$$

where  $E$  is the expectation operator, given  $g(x)$ .

# Monte Carlo Integration: Simple Example

- Previous result: can express an integral as an expectation relative to a (arbitrary, subject to obvious regularity conditions) density function.
- Use the law of large numbers (LLN) to approximate the expectation.
  - step 1: draw  $x_i$  independently from density,  $g$ , for  $i = 1, \dots, M$ .
  - step 2: evaluate  $f(x_i) / g(x_i)$  and compute:

$$\mu_M \equiv \frac{1}{M} \sum_{i=1}^M \frac{f(x_i)}{g(x_i)} \xrightarrow{M \rightarrow \infty} E \frac{f(x)}{g(x)}.$$

- Exercise.
  - Consider an integral where you have an analytic solution available, e.g.,  $\int_0^1 x^2 dx$ .
  - Evaluate the accuracy of the Monte Carlo method using various distributions on  $[0, 1]$  like uniform or Beta.

# Monte Carlo Integration: Simple Example

- Standard classical sampling theory applies.
- Independence of  $f(x_i) / g(x_i)$  over  $i$  implies:

$$\text{var} \left( \frac{1}{M} \sum_{i=1}^M \frac{f(x_i)}{g(x_i)} \right) = \frac{v_M}{M},$$

$$v_M \equiv \text{var} \left( \frac{f(x_i)}{g(x_i)} \right) \simeq \frac{1}{M} \sum_{i=1}^M \left[ \frac{f(x_i)}{g(x_i)} - \mu_M \right]^2.$$

- Central Limit Theorem

- Estimate of  $\int_a^b f(x) dx$  is a realization from a Normal distribution with mean estimated by  $\mu_M$  and variance,  $v_M/M$ .
- With 95% probability,

$$\mu_M - 1.96 \times \sqrt{\frac{v_M}{M}} \leq \int_a^b f(x) dx \leq \mu_M + 1.96 \times \sqrt{\frac{v_M}{M}}$$

- Pick  $g$  to minimize variance in  $f(x_i) / g(x_i)$  and  $M$  to minimize (subject to computing cost)  $v_M/M$ .

# Markov Chain, Monte Carlo (MCMC) Algorithms

- Among the top 10 algorithms "with the greatest influence on the development and practice of science and engineering in the 20th century".
  - Reference: January/February 2000 issue of Computing in Science & Engineering, a joint publication of the American Institute of Physics and the IEEE Computer Society.'
- Developed in 1946 by John von Neumann, Stan Ulam, and Nick Metropolis (see <http://www.siam.org/pdf/news/637.pdf>)



# MCMC Algorithm: Overview

- compute a sequence,  $\theta^{(1)}, \theta^{(2)}, \dots, \theta^{(M)}$ , of values of the  $N \times 1$  vector of model parameters in such a way that

$$\lim_{M \rightarrow \infty} \text{Frequency} \left[ \theta^{(i)} \text{ close to } \theta \right] = p \left( \theta | Y^{data} \right).$$

- Use  $\theta^{(1)}, \theta^{(2)}, \dots, \theta^{(M)}$  to obtain an approximation for
  - $E\theta, \text{Var}(\theta)$  under posterior distribution,  $p(\theta | Y^{data})$
  - $g(\theta^i | Y^{data}) = \int_{\theta_{i \neq j}} p(\theta | Y^{data}) d\theta d\theta$
  - $p(Y^{data}) = \int_{\theta} p(Y^{data} | \theta) p(\theta) d\theta$
  - posterior distribution of any function of  $\theta, f(\theta)$  (e.g., impulse responses functions, second moments).
- MCMC also useful for computing posterior mode,  $\arg \max_{\theta} p(\theta | Y^{data})$ .

# MCMC Algorithm: setting up

- Let  $G(\theta)$  denote the log of the posterior distribution (excluding an additive constant):

$$G(\theta) = \log p(Y^{data}|\theta) + \log p(\theta);$$

- Compute posterior mode:

$$\theta^* = \arg \max_{\theta} G(\theta).$$

- Compute the positive definite matrix,  $V$  :

$$V \equiv \left[ -\frac{\partial^2 G(\theta)}{\partial \theta \partial \theta'} \right]_{\theta=\theta^*}^{-1}$$

- Later, we will see that  $V$  is a rough estimate of the variance-covariance matrix of  $\theta$  under the posterior distribution.

# MCMC Algorithm: Metropolis-Hastings

- $\theta^{(1)} = \theta^*$
- to compute  $\theta^{(r)}$ , for  $r > 1$ 
  - step 1: select candidate  $\theta^{(r)}$ ,  $x$ ,

draw  $\underbrace{x}_{N \times 1}$  from  $\theta^{(r-1)} + \overbrace{k \times N}^{\text{'jump' distribution}} \left( \underbrace{0}_{N \times 1}, V \right)$ ,  $k$  is a scalar

- step 2: compute scalar,  $\lambda$  :

$$\lambda = \frac{p(Y^{data}|x) p(x)}{p(Y^{data}|\theta^{(r-1)}) p(\theta^{(r-1)})}$$

- step 3: compute  $\theta^{(r)}$  :

$$\theta^{(r)} = \begin{cases} \theta^{(r-1)} & \text{if } u > \lambda \\ x & \text{if } u < \lambda \end{cases}, \quad u \text{ is a realization from uniform } [0, 1]$$

# Practical issues

- What is a sensible value for  $k$ ?
  - set  $k$  so that you accept (i.e.,  $\theta^{(r)} = x$ ) in step 3 of MCMC algorithm are roughly 23 percent of time
- What value of  $M$  should you set?
  - want ‘convergence’, in the sense that if  $M$  is increased further, the econometric results do not change substantially
  - in practice,  $M = 10,000$  (a small value) up to  $M = 1,000,000$ .
  - large  $M$  is time-consuming.
    - could use Laplace approximation (after checking its accuracy) in initial phases of research project.
    - more on Laplace below.
- Burn-in: in practice, some initial  $\theta^{(i)}$ 's are discarded to minimize the impact of initial conditions on the results.
- Multiple chains: may promote efficiency.
  - increase independence among  $\theta^{(i)}$ 's.
  - can do MCMC utilizing parallel computing (Dynare can do this).

# MCMC Algorithm: Why Does it Work?

- Proposition that MCMC works may be surprising.
  - Whether or not it works does *not* depend on the details, i.e., precisely how you choose the jump distribution (of course, you had better use  $k > 0$  and  $V$  positive definite).
    - Proof: see, e.g., Robert, C. P. (2001), *The Bayesian Choice*, Second Edition, New York: Springer-Verlag.
  - The details may matter by improving the efficiency of the MCMC algorithm, i.e., by influencing what value of  $M$  you need.
- Some Intuition
  - the sequence,  $\theta^{(1)}, \theta^{(2)}, \dots, \theta^{(M)}$ , is relatively heavily populated by  $\theta$ 's that have high probability and relatively lightly populated by low probability  $\theta$ 's.
  - Additional intuition can be obtained by positing a simple scalar distribution and using MATLAB to verify that MCMC approximates it well (see, e.g., question 2 in assignment 9).

# MCMC Algorithm: using the Results

- To approximate marginal posterior distribution,  $g(\theta_i | Y^{data})$ , of  $\theta_i$ ,
  - compute and display the histogram of  $\theta_i^{(1)}, \theta_i^{(2)}, \dots, \theta_i^{(M)}$ ,  $i = 1, \dots, M$ .
- Other objects of interest:
  - mean and variance of posterior distribution  $\theta$  :

$$E\theta \simeq \bar{\theta} \equiv \frac{1}{M} \sum_{j=1}^M \theta^{(j)}, \quad \text{Var}(\theta) \simeq \frac{1}{M} \sum_{j=1}^M [\theta^{(j)} - \bar{\theta}] [\theta^{(j)} - \bar{\theta}]'.$$

# MCMC Algorithm: using the Results

- More complicated objects of interest:
  - impulse response functions,
  - model second moments,
  - forecasts,
  - Kalman smoothed estimates of real rate, natural rate, etc.
- All these things can be represented as non-linear functions of the model parameters, i.e.,  $f(\theta)$ .
  - can approximate the distribution of  $f(\theta)$  using

$$f(\theta^{(1)}), \dots, f(\theta^{(M)})$$

$$\rightarrow Ef(\theta) \simeq \bar{f} \equiv \frac{1}{M} \sum_{i=1}^M f(\theta^{(i)}),$$

$$Var(f(\theta)) \simeq \frac{1}{M} \sum_{i=1}^M [f(\theta^{(i)}) - \bar{f}] [f(\theta^{(i)}) - \bar{f}]'$$

# MCMC: Remaining Issues

- In addition to the first and second moments already discussed, would also like to have the marginal likelihood of the data.
- Marginal likelihood is a Bayesian measure of model fit.



# MCMC Algorithm: the Marginal Likelihood

- Consider the following sample average:

$$\frac{1}{M} \sum_{j=1}^M \frac{h(\theta^{(j)})}{p(Y^{data}|\theta^{(j)}) p(\theta^{(j)})},$$

where  $h(\theta)$  is an arbitrary density function over the  $N$ -dimensional variable,  $\theta$ .

By the law of large numbers,

$$\frac{1}{M} \sum_{j=1}^M \frac{h(\theta^{(j)})}{p(Y^{data}|\theta^{(j)}) p(\theta^{(j)})} \xrightarrow{M \rightarrow \infty} E \left( \frac{h(\theta)}{p(Y^{data}|\theta) p(\theta)} \right)$$

# MCMC Algorithm: the Marginal Likelihood

$$\begin{aligned} \frac{1}{M} \sum_{j=1}^M \frac{h(\theta^{(j)})}{p(Y^{data}|\theta^{(j)}) p(\theta^{(j)})} &\rightarrow_{M \rightarrow \infty} E \left( \frac{h(\theta)}{p(Y^{data}|\theta) p(\theta)} \right) \\ &= \int_{\theta} \left( \frac{h(\theta)}{p(Y^{data}|\theta) p(\theta)} \right) \frac{p(Y^{data}|\theta) p(\theta)}{p(Y^{data})} d\theta = \frac{1}{p(Y^{data})}. \end{aligned}$$

- When  $h(\theta) = p(\theta)$ , *harmonic mean estimator of the marginal likelihood*.
- Ideally, want an  $h$  such that the variance of

$$\frac{h(\theta^{(j)})}{p(Y^{data}|\theta^{(j)}) p(\theta^{(j)})}$$

is small (recall the earlier discussion of Monte Carlo integration). More on this below.

# Laplace Approximation to Posterior Distribution

- In practice, MCMC algorithm very time intensive.
- Laplace approximation is easy to compute and in many cases it provides a 'quick and dirty' approximation that is quite good.

Let  $\theta \in R^N$  denote the  $N$ -dimensional vector of parameters and, as before,

$$G(\theta) \equiv \log p(Y^{data}|\theta) p(\theta)$$

$$p(Y^{data}|\theta) \sim \text{likelihood of data}$$

$$p(\theta) \sim \text{prior on parameters}$$

$$\theta^* \sim \text{maximum of } G(\theta) \text{ (i.e., mode)}$$

# Laplace Approximation

Second order Taylor series expansion of  
 $G(\theta) \equiv \log [p(Y^{data}|\theta) p(\theta)]$  about  $\theta = \theta^*$  :

$$G(\theta) \approx G(\theta^*) + G_{\theta}(\theta^*)(\theta - \theta^*) - \frac{1}{2}(\theta - \theta^*)' G_{\theta\theta}(\theta^*)(\theta - \theta^*),$$

where

$$G_{\theta\theta}(\theta^*) = - \frac{\partial^2 \log p(Y^{data}|\theta) p(\theta)}{\partial\theta\partial\theta'} \Big|_{\theta=\theta^*}$$

Interior optimality of  $\theta^*$  implies:

$$G_{\theta}(\theta^*) = 0, \quad G_{\theta\theta}(\theta^*) \text{ positive definite}$$

Then:

$$\begin{aligned} & p(Y^{data}|\theta) p(\theta) \\ \simeq & p(Y^{data}|\theta^*) p(\theta^*) \exp \left\{ -\frac{1}{2}(\theta - \theta^*)' G_{\theta\theta}(\theta^*)(\theta - \theta^*) \right\}. \end{aligned}$$

# Laplace Approximation to Posterior Distribution

Property of Normal distribution:

$$\int_{\theta} \frac{1}{(2\pi)^{\frac{N}{2}}} |G_{\theta\theta}(\theta^*)|^{\frac{1}{2}} \exp \left\{ -\frac{1}{2} (\theta - \theta^*)' G_{\theta\theta}(\theta^*) (\theta - \theta^*) \right\} d\theta = 1$$

Then,

$$\begin{aligned} \int p(Y^{data}|\theta) p(\theta) d\theta &\simeq \int p(Y^{data}|\theta^*) p(\theta^*) \\ &\quad \times \exp \left\{ -\frac{1}{2} (\theta - \theta^*)' G_{\theta\theta}(\theta^*) (\theta - \theta^*) \right\} \\ &= \frac{p(Y^{data}|\theta^*) p(\theta^*)}{\frac{1}{(2\pi)^{\frac{N}{2}}} |G_{\theta\theta}(\theta^*)|^{\frac{1}{2}}}. \end{aligned}$$

# Laplace Approximation

- Conclude:

$$p(Y^{data}) \simeq \frac{p(Y^{data}|\theta^*) p(\theta^*)}{\frac{1}{(2\pi)^{\frac{N}{2}}} |G_{\theta\theta}(\theta^*)|^{\frac{1}{2}}}.$$

- Laplace approximation to posterior distribution:

$$\frac{p(Y^{data}|\theta) p(\theta)}{p(Y^{data})} \simeq \frac{1}{(2\pi)^{\frac{N}{2}}} |G_{\theta\theta}(\theta^*)|^{\frac{1}{2}} \times \exp\left\{-\frac{1}{2}(\theta - \theta^*)' G_{\theta\theta}(\theta^*) (\theta - \theta^*)\right\}$$

- So, posterior of  $\theta_i$  (i.e.,  $g(\theta_i|Y^{data})$ ) is approximately

$$\theta_i \sim N\left(\theta_i^*, \left[G_{\theta\theta}(\theta^*)^{-1}\right]_{ii}\right).$$

# Modified Harmonic Mean Estimator of Marginal Likelihood

- Harmonic mean estimator of the marginal likelihood,  $p(Y^{data})$ :

$$\left[ \frac{1}{M} \sum_{j=1}^M \frac{h(\theta^{(j)})}{p(Y^{data}|\theta^{(j)}) p(\theta^{(j)})} \right]^{-1},$$

with  $h(\theta)$  set to  $p(\theta)$ .

- In this case, the marginal likelihood is the harmonic mean of the likelihood, evaluated at the values of  $\theta$  generated by the MCMC algorithm.
- Problem: the variance of the object being averaged is likely to be high, requiring high  $M$  for accuracy.
- When  $h(\theta)$  is instead equated to Laplace approximation of posterior distribution, then  $h(\theta)$  is approximately proportional to  $p(Y^{data}|\theta^{(j)}) p(\theta^{(j)})$  so that the variance of the variable being averaged in the last expression is low.

# The Marginal Likelihood and Model Comparison

- Suppose we have two models, *Model 1* and *Model 2*.
  - compute  $p(Y^{data}|Model\ 1)$  and  $p(Y^{data}|Model\ 2)$
- Suppose  $p(Y^{data}|Model\ 1) > p(Y^{data}|Model\ 2)$ . Then, posterior odds on Model 1 higher than Model 2.
  - ‘Model 1 fits better than Model 2’
- Can use this to compare across two different models, or to evaluate contribution to fit of various model features: habit persistence, adjustment costs, etc.
  - For an application of this and the other methods in these notes, see Smets and Wouters, AER 2007.