

Simulation-Based Approaches to Bayesian Inference

Frank Schorfheide

Department of Economics
University of Pennsylvania

Introduction

- For more complicated models it is often infeasible to derive moments of the posterior distribution analytically and to calculate the marginal data density.
- Instead: use algorithm that enables us to generate draws from posterior distribution.
- Based on draws, calculate numerical approximations to moments of interest.
- Important algorithms: Gibbs sampler and Metropolis-Hastings:
 - Gibbs sampler: for instance, for VARs
 - Metropolis-Hastings: for instance, DSGE models
- Markov-Chain Monte Carlo (MCMC) algorithms: construct a Markov Chain with ergodic distribution $p(\theta|Y)$.

Numerical Integration

- The posterior expectation of a function $h(\theta)$ is equal to the integral

$$\mathbb{E}[h(\theta)|Y] = \int h(\theta)p(\theta|Y)d\theta \quad (1)$$

- **MC Integration:**

(i) Requirements: It is possible to generate draws from the posterior density $p(\theta|Y)$.

(ii) Steps:

1. For $s = 1, \dots, n_{sim}$ draw $\theta^{(s)}$ from the posterior density $p(\theta|Y)$.
2. Approximate the integral by

$$\mathbb{E}[h(\theta)|Y] \approx \frac{1}{n_{sim}} \sum_{s=1}^{n_{sim}} h(\theta^{(s)}) \quad \square \quad (2)$$

- Use Law of Large Numbers for $h(\theta^{(s)})$ and Central Limit Theorem to assess the approximation error.

Importance Sampling

- Suppose it is not possible to generate draws from $p(\theta|Y)$ on the computer, but one can calculate the posterior density up to a constant, that is

$$q(\theta|Y) = cp(\theta|Y) \tag{3}$$

and generate draws from a normalized density $g(\theta)$.

- Note that

$$\begin{aligned} \mathbb{E}[h(\theta)|Y] &= \int h(\theta)p(\theta|Y)d\theta \\ &= \int h(\theta)\frac{p(\theta|Y)}{g(\theta)}g(\theta)d\theta \\ &= \frac{1}{c} \int h(\theta)\frac{q(\theta|Y)}{g(\theta)}g(\theta)d\theta \end{aligned} \tag{4}$$

Importance Sampling

- Moreover,

$$\int \frac{q(\theta|Y)}{g(\theta)} g(\theta) d\theta = c \int p(\theta|Y) d\theta = c \quad (5)$$

- Define the importance ratios

$$w(\theta) = q(\theta|Y)/g(\theta) \quad (6)$$

- The above analysis suggests the following algorithm, known as importance sampling...

Importance Sampling

- **MC Integration through Importance Sampling:**

(i) Requirements: Posterior density $p(\theta|Y)$ can be evaluated up to a constant: $q(\theta|Y) = cp(\theta|Y)$. It is possible to generate draws from a normalized density $g(\theta)$. The draws need not be independent.

(ii) Steps:

1. For $s = 1, \dots, n_{sim}$ draw $\theta^{(s)}$ from the density $g(\theta)$.

2. Approximate the integral by

$$\mathbb{E}[h(\theta)|Y] \approx \frac{\frac{1}{n_{sim}} \sum_{s=1}^{n_{sim}} w(\theta^{(s)})h(\theta^{(s)})}{\frac{1}{n_{sim}} \sum_{s=1}^{n_{sim}} w(\theta^{(s)})} \quad \square \quad (7)$$

Importance Sampling

- In principle the approximation error converges to zero for a large class of densities $g(\theta)$.
- The important question is: how fast?
- The algorithm will work well if $g(\theta)$ can be chosen such that

$$h(\theta)w(\theta)$$

stays roughly constant across draws $\theta^{(s)}$.

- The convergence will be very slow if the importance ratios are very small with high probability and very large with low probability.
- Consider CLT: If $\{y_t\}$ is strictly stationary and ergodic with $\mathbb{E}[y_1] = 0$, $\mathbb{E}[y_1^2] = \sigma^2 < \infty$, and $\bar{\sigma}_T = \text{var}(T^{-1/2} \sum y_t) \longrightarrow \bar{\sigma}^2 < \infty$, then

$$\frac{1}{\sqrt{T\bar{\sigma}_T}} \sum y_t \xrightarrow{p} \mathcal{N}(0, 1). \quad \square$$

Data Augmentation Algorithms

- The data augmentation algorithm is based on two simple identities.
- The *posterior identity* is of the form

$$p(\theta|Y) = \int_Z p(\theta|Y, Z)p(Z|Y)dZ \quad (8)$$

- and the *predictive identity* is

$$p(Z|Y) \int p(Z|\theta, Y)p(\theta|Y)d\theta \quad (9)$$

- When substituting the latter identity into the former we obtain

$$\begin{aligned} p(\theta|Y) &= \int_Z p(\theta|Y, Z) \left[\int_{\tilde{\theta}} p(Z|\tilde{\theta}, Y)p(\tilde{\theta}|Y)d\tilde{\theta} \right] dZ \\ &= \int_{\tilde{\theta}} \left[\int_Z p(\theta|Y, Z)p(Z|Y, \tilde{\theta})dZ \right] p(\tilde{\theta}|Y)d\tilde{\theta} \\ &= \int_{\tilde{\theta}} K(\theta|\tilde{\theta}, Y)p(\tilde{\theta}|Y)d\tilde{\theta} \end{aligned} \quad (10)$$

Data Augmentation Algorithms

- Thus, $p(\theta|Y)$ has to satisfy an integral equation. Let $g(\theta)$ be a normalized probability density. Define the mapping

$$M[g(\theta)] = \int K(\theta|\tilde{\theta}, Y)g(\tilde{\theta})d\tilde{\theta} \quad (11)$$

M maps a density $g(\theta)$ into a density $g'(\theta)$. Let $g_{(s+1)}(\theta) = M[g_{(s)}(\theta)]$.

- Under suitable regularity conditions (see for instance Tanner, 1996)
 - (i) The fixed point $g^*(\theta)$ of the mapping $M[g(\theta)] = g(\theta)$ is unique.
 - (ii) The mapping M is a contraction mapping and the sequence of densities $\{g_{(s)}(\theta)\}$ converges to the fixed point $g^*(\theta)$

$$\int |g_{(s)}(\theta) - g^*(\theta)|d\theta \longrightarrow 0$$

as $s \longrightarrow \infty$. \square

Data Augmentation Algorithms

- Thus, if one applies the mapping M to some initial density $g_{(0)}(\theta)$ one converges eventually to the solution of the integral equation.

- **Data Augmentation Algorithm**

(i) Requirements: It is easily possible to generate draws from $p(Z|\theta, Y)$ and $p(\theta|Z, Y)$.

(ii) The following two steps are repeated for $s = 1, \dots, n_{sim}$.

1. Generate a sample Z_1, \dots, Z_m from the current approximation to $p(Z|Y)$ as follows: For $i = 1, \dots, m$ draw $\theta_{(i)}$ from $p_{(s)}(\theta|Y)$ and draw $Z_{(i)}$ from $p(Z|\theta_{(i)}, Y)$.
2. Use the posterior identity to update from $p_{(s)}(\theta|Y)$ to $p_{(s+1)}(\theta|Y)$:

$$p_{(s+1)}(\theta|Y) = \frac{1}{m} \sum_{i=1}^m p(\theta|Z_{(i)}, Y)$$

A draw of θ can be easily generated by drawing an $i = 1, \dots, m$ with uniform probability $1/m$ and a θ from $p(\theta|Z_{(i)}, Y)$

Data Augmentation Algorithms

- For large m and large n_{sim} the mixture $p_{(s+1)}(\theta|Y)$ will provide a close approximation of $p(\theta|Y)$. In many cases we are not so much interested in the density $p(\theta|Y)$ but rather in random draws from this density. Choosing $m = 1$ leads to *chained data augmentation*.
- **Chained Data Augmentation Algorithm**
 - (i) Requirements: It is easily possible to generate draws from $p(Z|\theta, Y)$ and $p(\theta|Z, Y)$.
 - (ii) The following two steps are repeated for $s = 1, \dots, n_{sim}$.
 1. Draw $Z^{(s+1)}$ from the density $p(Z|\theta^{(s)}, Y)$.
 2. Draw $\theta^{(s)}$ from the density $p(\theta|Z^{(s+1)}, Y)$. \square

Data Augmentation Algorithms

- The sequence $\{\theta^{(s)}\}_{s=1}^{n_{sim}}$ in the chained data augmentation forms a Markov Chain with transition Kernel

$$K(\theta|\tilde{\theta}, Y) = \int p(\theta|Y, Z)p(Z|Y, \tilde{\theta})dZ \quad (12)$$

- The stationary distribution of this Markov chain has to satisfy the integral equation

$$p(\theta|Y) = \int K(\theta|\tilde{\theta}, Y)p(\tilde{\theta}|Y)d\tilde{\theta} \quad (13)$$

- Hence for large s , $\{\theta^{(s)}\}_{s=1}^{n_{sim}}$ are draws from the posterior distribution $p(\theta|Y)$. Moreover, the Z 's obtained from the chained data augmentation algorithm are draws from $p(Z|Y)$. This Fact is exploited by the Gibbs Sampler.

Data Augmentation Algorithms

- **Gibbs Sampler**

(i) Requirements: Suppose the parameter vector θ can be partitioned into $\theta = [\theta'_1, \dots, \theta'_m]'$.

For each i it is possible to generate draws of θ_i from the conditional distribution

$p(\theta_i | \theta_{-i}, Y)$ where θ_{-i} denotes the vector θ without the partition θ_i .

(ii) The following steps are repeated for $s = 1, \dots, n_{sim}$.

- Draw $\theta_1^{(s+1)}$ from the density $p(\theta_1 | \theta_2^{(s)}, \dots, \theta_m^{(s)}, Y)$.

- Draw $\theta_2^{(s+1)}$ from the density $p(\theta_2 | \theta_1^{(s+1)}, \theta_3^{(s)}, \dots, \theta_m^{(s)}, Y)$.

- ...

- Draw $\theta_m^{(s+1)}$ from the density $p(\theta_m | \theta_1^{(s+1)}, \dots, \theta_{m-1}^{(s+1)}, Y)$. \square

Data Augmentation Algorithms

- For large s we obtain dependent draws from the posterior distribution of θ . It is common practice to discard the initial draws.
- Approximate the mean and covariance matrix of θ by Monte Carlo averages:

$$\widehat{E}[\theta] = \frac{1}{n_{sim} - n_0} \sum_{s=n_0+1}^{n_{sim}} \theta^{(s)}.$$

Example: Gibbs Sampler

- Simple example: suppose our posterior is of the form

$$\begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \tilde{\theta}_1 \\ \tilde{\theta}_2 \end{bmatrix}, \begin{bmatrix} \tilde{V}_{11} & \tilde{V}_{12} \\ & \tilde{V}_{22} \end{bmatrix} \right).$$

- In practice the posterior mean $\tilde{\theta}$ and variance \tilde{V} are functions of data and prior.

- Let's assume:

$$\begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} 0.5 \\ -0.5 \end{bmatrix}, \begin{bmatrix} 1 & 0.5 \\ & 1 \end{bmatrix} \right).$$

Example: Gibbs Sampler

- In our example the conditional distributions are...
- Distribution of $\theta_1|\theta_2$:

$$\mathcal{N}\left(\mu_1 + \sigma_{12}\sigma_{22}^{-1}(\theta_2 - \mu_2), \sigma_{11} - \sigma_{12}\sigma_{22}^{-1}\sigma_{21}\right)$$


- Distribution of $\theta_2|\theta_1$:

$$\mathcal{N}\left(\mu_2 + \sigma_{21}\sigma_{11}^{-1}(\theta_1 - \mu_1), \sigma_{22} - \sigma_{21}\sigma_{11}^{-1}\sigma_{12}\right)$$


- We consider cumulative means of θ draws... ($\theta_2^{(0)} = 10$)
- and autocorrelation function of θ draws.

(Figures)


$E(\theta_1)$




$E(\theta_2)$




$VAR(\theta_1)$




$COV(\theta_1, \theta_2)$




$E(\theta_1)$




$E(\theta_2)$



$\text{VAR}(\theta_1)$



$\text{COV}(\theta_1, \theta_2)$



Metropolis-Hastings Algorithm

- Consider an m -state Markov process x_t .
- The possible states are denoted by $S = \{s_1, \dots, s_m\}$.
- Transition probabilities:

$$P = \begin{pmatrix} p_{11} & \cdots & p_{1m} \\ \vdots & \ddots & \vdots \\ p_{m1} & \cdots & p_{mm} \end{pmatrix} \quad (14)$$

where p_{ij} is the probability of moving from state i to state j .

- Let $w(t) = [w_1(t), \dots, w_m(t)]$ be a $1 \times m$ vector of probabilities of x_t being in state i in period t .
- The corresponding probabilities for period $t + 1$ are $w(t + 1) = w(t)P$.

Metropolis-Hastings Algorithm

- π is equilibrium distribution if $\pi = \pi P$.
- A Markov chain is reversible if probability of $i \mapsto j$ is the same as $j \mapsto i$:

$$\pi_i p_{ij} = \pi_j p_{ji}$$

- A chain that is reversible has an equilibrium distribution π because

$$(\pi P)_j = \sum_{i=1}^m \pi_i p_{ij} = \sum_{i=1}^m \pi_j p_{ji} = \pi_j \sum_{i=1}^m p_{ji} = \pi_j \quad (15)$$

- To sample from the equilibrium distribution, one can start the chain from any $w(0)$ until it settles down to the equilibrium distribution.

Metropolis-Hastings Algorithm

- Idea of Metropolis algorithm: construct transition matrix P from a transition matrix Q such that P has desired equilibrium distribution π .
- Why: we don't know how to draw from the posterior $p(\theta|Y^T)$ (corresponds to π in our example), but we know how to draw from a normal distribution (corresponds to using Q).
- Time t iteration: suppose we are in state s_i Based on Q draw a proposed state s_j .
 - With probability α_{ij} proposed state is accepted. Move from s_i to s_j .
 - With probability $1 - \alpha_{ij}$ stay in state s_i .
- Choice of α_{ij} ?

Metropolis-Hastings Algorithm

- Choose $\alpha_{ij} = \min [1, \pi_j/\pi_i]$.
- The resulting chain is reversible and has equilibrium distribution π :

$$\begin{aligned}\pi_i p_{ij} &= \pi_i \min[1, \pi_j/\pi_i] q_{ij} \\ &= \min[\pi_i, \pi_j] q_{ij} \\ &= \min[\pi_i, \pi_j] q_{ji} \\ &= \pi_j p_{ji}\end{aligned}\tag{16}$$

- Weak regularity conditions on Q can ensure that the equilibrium distribution π is unique and that the chain is convergent.
- These ideas can be generalized to the continuous case.

Metropolis-Hastings Algorithm

- Special Case: Random Walk Metropolis Algorithm
 - Initialization: Choose a $\theta^{(0)}$ to initialize the chain.
 - Step s_1 : Draw candidate parameter vector ψ from a jumping distribution $J_s(\psi|\theta^{(s-1)})$.

Example: draw ψ from the distribution

$$\psi \sim \mathcal{N}(\theta^{(s-1)}, c^2 I)$$

where I is the identity matrix and c is a scalar tuning parameter for the algorithm.

- Step s_2 :

$$\theta^{(s)} = \begin{cases} \psi & \text{with probability } \min \left\{ 1, \frac{p(\psi|Y)}{p(\theta^{(s-1)}|Y)} \right\} \\ \theta^{(s-1)} & \text{otherwise} \end{cases}$$

We refer to $\theta^{(s)} = \theta^{(s-1)}$ as “rejection” of the proposed step.

Execute steps s_1 and s_2 for $s = 1, \dots, n_{sim}$.

Metropolis-Hastings Algorithm

- We saw: the sequence of draws $\{\theta^{(s)}\}$ is serially correlated.
- It typically satisfies a weak law of large numbers, that is,

$$\hat{\mathbb{E}}[\theta|Y^T] = \frac{1}{n_s} \sum_{s=1}^{n_s} \theta^{(s)} \xrightarrow{p} \int \theta p(\theta|Y^T) d\theta. \quad (17)$$

- Under some regularity conditions it also satisfies a central limit theorem, meaning: Newey-West standard errors could be used to get numerical standard errors.
- Of course we can transform $\theta^{(s)}$ and evaluate:

$$\hat{\mathbb{E}}[g(\theta)|Y^T] = \frac{1}{n_s} \sum_{s=1}^{n_s} g(\theta^{(s)}) \xrightarrow{p?} \int g(\theta) p(\theta|Y^T) d\theta. \quad (18)$$

Use for: standard deviations, impulse response functions, etc...

- Univariate confidence intervals (connected): (i) sort draws, (ii) search for shortest interval that contains a pre-specified fraction, say 95 percent, of the draws.

Convergence?

- Issue: speed of convergence of Markov Chain and Monte Carlo averages.
- Informal assessment of convergence: plot

$$\frac{1}{n_s} \sum_{s=1}^{n_s} g(\theta^{(s)}) \quad (19)$$

as a function of n_s .

- Start Markov-Chain at over-dispersed (i.e., extreme) values of θ and check whether different runs of the chain settle to the same distribution.
- Huge literature, here some references: Gelman, Carlin, Stern and Rubin (1995), Tanner (1996), Geweke (1999).

Example: Metropolis-Hastings Algorithm

- Simple example: suppose our posterior is a mixture of normals of the form

$$\begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} \sim \left\{ \begin{array}{l} \mathcal{N} \left(\begin{bmatrix} \tilde{\theta}_{11} \\ \tilde{\theta}_{12} \end{bmatrix}, \begin{bmatrix} \tilde{V}_{11} & \tilde{V}_{12} \\ & \tilde{V}_{22} \end{bmatrix} \right) \text{ with probability } \frac{1}{2} \\ \mathcal{N} \left(\begin{bmatrix} \tilde{\theta}_{21} \\ \tilde{\theta}_{22} \end{bmatrix}, \begin{bmatrix} \tilde{V}_{11} & \tilde{V}_{12} \\ & \tilde{V}_{22} \end{bmatrix} \right) \text{ with probability } \frac{1}{2} \end{array} \right.$$

Example: Metropolis-Hastings Algorithm


- Simple example: suppose our posterior is a mixture of normals of the form

$$\begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} \sim \begin{cases} \mathcal{N} \left(\begin{bmatrix} 1.5 \\ 1.5 \end{bmatrix}, \begin{bmatrix} 1 & 0.5 \\ & 1 \end{bmatrix} \right) & \text{with probability } \frac{1}{2} \\ \mathcal{N} \left(\begin{bmatrix} -1.5 \\ -1.5 \end{bmatrix}, \begin{bmatrix} 1 & 0.5 \\ & 1 \end{bmatrix} \right) & \text{with probability } \frac{1}{2} \end{cases}$$


- Configuration of Algorithm

- Starting values: $\theta_1^{(1)} = 10$, $\theta_2^{(1)} = -10$
- Proposal density: $\mathcal{N}(\theta^{(s-1)}, 0.1^2 * \mathcal{I})$
- Number of draws: 2000


- Small steps... Rejection rate: 8.45 %.




THETA1 Draws




THETA2 Draws




$E(\theta_1)$




$E(\theta_2)$




$\text{VAR}(\theta_2)$




$\text{COV}(\theta_1, \theta_2)$




Theta 1




Theta 2



(Theta 1)²



(Theta 2)²



Example: Metropolis-Hastings Algorithm


- Simple example: suppose our posterior is a mixture of normals of the form

$$\begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} \sim \begin{cases} \mathcal{N} \left(\begin{bmatrix} 1.5 \\ 1.5 \end{bmatrix}, \begin{bmatrix} 1 & 0.5 \\ & 1 \end{bmatrix} \right) & \text{with probability } \frac{1}{2} \\ \mathcal{N} \left(\begin{bmatrix} -1.5 \\ -1.5 \end{bmatrix}, \begin{bmatrix} 1 & 0.5 \\ & 1 \end{bmatrix} \right) & \text{with probability } \frac{1}{2} \end{cases}$$


- Configuration of Algorithm

- Starting values: $\theta_1^{(1)} = 10$, $\theta_2^{(1)} = -10$
- Proposal density: $\mathcal{N}(\theta^{(s-1)}, 1^2 * \mathcal{I})$
- Number of draws: 2000


- Mid-size steps... Rejection rate: 45.25 %.




THETA1 Draws




THETA2 Draws




$E(\theta_1)$




$E(\theta_2)$



$\text{VAR}(\theta_2)$



$\text{COV}(\theta_1, \theta_2)$



Example: Metropolis-Hastings Algorithm

- Simple example: suppose our posterior is a mixture of normals of the form

$$\begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} \sim \begin{cases} \mathcal{N} \left(\begin{bmatrix} 1.5 \\ 1.5 \end{bmatrix}, \begin{bmatrix} 1 & 0.5 \\ & 1 \end{bmatrix} \right) & \text{with probability } \frac{1}{2} \\ \mathcal{N} \left(\begin{bmatrix} -1.5 \\ -1.5 \end{bmatrix}, \begin{bmatrix} 1 & 0.5 \\ & 1 \end{bmatrix} \right) & \text{with probability } \frac{1}{2} \end{cases}$$


- Configuration of Algorithm

– Starting values: $\theta_1^{(1)} = 10$, $\theta_2^{(1)} = -10$


– Proposal density: $\mathcal{N}(\theta^{(s-1)}, 8^2 * \mathcal{I})$

– Number of draws: 2000


- Large steps... Rejection rate: 95.50 %.




THETA1 Draws




THETA2 Draws




$E(\theta_1)$




$E(\theta_2)$



$\text{VAR}(\theta_2)$



$\text{COV}(\theta_1, \theta_2)$



Example: Metropolis-Hastings Algorithm

- Move modes of mixture further apart...

$$\begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} \sim \begin{cases} \mathcal{N} \left(\begin{bmatrix} 3 \\ 3 \end{bmatrix}, \begin{bmatrix} 1 & 0.5 \\ & 1 \end{bmatrix} \right) & \text{with probability } \frac{1}{2} \\ \mathcal{N} \left(\begin{bmatrix} -3 \\ -3 \end{bmatrix}, \begin{bmatrix} 1 & 0.5 \\ & 1 \end{bmatrix} \right) & \text{with probability } \frac{1}{2} \end{cases}$$


- Configuration of Algorithm

– Starting values: $\theta_1^{(1)} = 10$, $\theta_2^{(1)} = -10$


– Proposal density: $\mathcal{N}(\theta^{(s-1)}, 1^2 * \mathcal{I})$

– Number of draws: 2000


- Rejection rate: 48.35 %.




THETA1 Draws




THETA2 Draws




$E(\theta_1)$




$E(\theta_2)$



$\text{VAR}(\theta_2)$



$\text{COV}(\theta_1, \theta_2)$



Example: Metropolis-Hastings Algorithm


- Simple example: increase number of draws...


$$\begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} \sim \begin{cases} \mathcal{N} \left(\begin{bmatrix} 3 \\ 3 \end{bmatrix}, \begin{bmatrix} 1 & 0.5 \\ & 1 \end{bmatrix} \right) & \text{with probability } \frac{1}{2} \\ \mathcal{N} \left(\begin{bmatrix} -3 \\ -3 \end{bmatrix}, \begin{bmatrix} 1 & 0.5 \\ & 1 \end{bmatrix} \right) & \text{with probability } \frac{1}{2} \end{cases}$$

- Configuration of Algorithm


- Starting values: $\theta_1^{(1)} = 10, \theta_2^{(1)} = 10$
- Proposal density: $\mathcal{N}(\theta^{(s-1)}, 1^2 * \mathcal{I})$
- Number of draws: 20000

- Rejection rate: 48.82 %.





THETA1 Draws



THETA2 Draws

