

The Metropolis-Hastings Algorithm

Frank Schorfheide
University of Pennsylvania
EABCN Training School

May 10, 2016

Markov Chain Monte Carlo (MCMC)

- Main idea: create a sequence of serially correlated draws such that the distribution of θ^i converges to the posterior distribution $p(\theta|Y)$.

Generic Metropolis-Hastings Algorithm

For $i = 1$ to N :

- 1 Draw ϑ from a density $q(\vartheta|\theta^{i-1})$.
- 2 Set $\theta^i = \vartheta$ with probability

$$\alpha(\vartheta|\theta^{i-1}) = \min \left\{ 1, \frac{p(Y|\vartheta)p(\vartheta)/q(\vartheta|\theta^{i-1})}{p(Y|\theta^{i-1})p(\theta^{i-1})/q(\theta^{i-1}|\vartheta)} \right\}$$

and $\theta^i = \theta^{i-1}$ otherwise.

Recall $p(\theta|Y) \propto p(Y|\theta)p(\theta)$.

We draw θ^i conditional on a parameter draw θ^{i-1} : leads to Markov transition kernel $K(\theta|\tilde{\theta})$.

- It can be shown that

$$p(\theta|Y) = \int K(\theta|\tilde{\theta})p(\tilde{\theta}|Y)d\tilde{\theta}.$$

- Write

$$K(\theta|\tilde{\theta}) = u(\theta|\tilde{\theta}) + r(\tilde{\theta})\delta_{\tilde{\theta}}(\theta).$$

- $u(\theta|\tilde{\theta})$ is the density kernel (note that $u(\theta|\cdot)$ does not integrate to one) for accepted draws:

$$u(\theta|\tilde{\theta}) = \alpha(\theta|\tilde{\theta})q(\theta|\tilde{\theta}).$$

- Rejection probability:

$$r(\tilde{\theta}) = \int [1 - \alpha(\theta|\tilde{\theta})]q(\theta|\tilde{\theta})d\theta = 1 - \int u(\theta|\tilde{\theta})d\theta.$$

Importance Invariance Property

- Reversibility: Conditional on the sampler not rejecting the proposed draw, the density associated with a transition from $\tilde{\theta}$ to θ is identical to the density associated with a transition from θ to $\tilde{\theta}$:

$$\begin{aligned} p(\tilde{\theta}|Y)u(\theta|\tilde{\theta}) &= p(\tilde{\theta}|Y)q(\theta|\tilde{\theta}) \min \left\{ 1, \frac{p(\theta|Y)/q(\theta|\tilde{\theta})}{p(\tilde{\theta}|Y)/q(\tilde{\theta}|\theta)} \right\} \\ &= \min \{ p(\tilde{\theta}|Y)q(\theta|\tilde{\theta}), p(\theta|Y)q(\tilde{\theta}|\theta) \} \\ &= p(\theta|Y)q(\tilde{\theta}|\theta) \min \left\{ \frac{p(\tilde{\theta}|Y)/q(\tilde{\theta}|\theta)}{p(\theta|Y)/q(\theta|\tilde{\theta})}, 1 \right\} \\ &= p(\theta|Y)u(\tilde{\theta}|\theta). \end{aligned}$$

- Using the reversibility result, we can now verify the invariance property:

$$\begin{aligned} \int K(\theta|\tilde{\theta})p(\tilde{\theta}|Y)d\tilde{\theta} &= \int u(\theta|\tilde{\theta})p(\tilde{\theta}|Y)d\tilde{\theta} + \int r(\tilde{\theta})\delta_{\tilde{\theta}}(\theta)p(\tilde{\theta}|Y)d\tilde{\theta} \\ &= \int u(\tilde{\theta}|\theta)p(\theta|Y)d\tilde{\theta} + r(\theta)p(\theta|Y) \\ &= p(\theta|Y) \end{aligned}$$

A Discrete Example

- Suppose parameter vector θ is scalar and takes only two values:

$$\Theta = \{\tau_1, \tau_2\}$$

- The posterior distribution $p(\theta|Y)$ can be represented by a set of probabilities collected in the vector π , say $\pi = [\pi_1, \pi_2]$ with $\pi_2 > \pi_1$.
- Suppose we obtain ϑ based on transition matrix Q :

$$Q = \begin{bmatrix} q & (1-q) \\ (1-q) & q \end{bmatrix}.$$

- Iteration i : suppose that $\theta^{i-1} = \tau_j$. Based on transition matrix

$$Q = \begin{bmatrix} q & (1-q) \\ (1-q) & q \end{bmatrix},$$

determine a proposed state $\vartheta = \tau_s$.

- With probability $\alpha(\tau_s|\tau_j)$ the proposed state is accepted. Set $\theta^i = \vartheta = \tau_s$.
 - With probability $1 - \alpha(\tau_s|\tau_j)$ stay in old state and set $\theta^i = \theta^{i-1} = \tau_j$.
- Choose (Q terms cancel because of symmetry)

$$\alpha(\tau_s|\tau_j) = \min \left\{ 1, \frac{\pi_s}{\pi_j} \right\}.$$

Discrete MH Algorithm: Transition Matrix

- The resulting chain's transition matrix is:

$$K = \begin{bmatrix} q & (1-q) \\ (1-q)\frac{\pi_1}{\pi_2} & q + (1-q)\left(1 - \frac{\pi_1}{\pi_2}\right) \end{bmatrix}.$$

- Straightforward calculations reveal that the transition matrix K has eigenvalues:

$$\lambda_1(K) = 1, \quad \lambda_2(K) = q - (1-q)\frac{\pi_1}{1 - \pi_1}.$$

- Equilibrium distribution is eigenvector associated with unit eigenvalue.
- For $q \in [0, 1)$ the equilibrium distribution is unique.

- The persistence of the Markov chain depends on second eigenvalue, which depends on the proposal distribution Q .
- Define the transformed parameter

$$\xi^i = \frac{\theta^i - \tau_1}{\tau_2 - \tau_1}.$$

- We can represent the Markov chain associated with ξ^i as first-order autoregressive process

$$\xi^i = (1 - k_{22}) + \lambda_2(K)\xi^{i-1} + \nu^i.$$

- Conditional on $\xi^i = j$, $j = 0, 1$, the innovation ν^i has support on k_{jj} and $(1 - k_{jj})$, its conditional mean is equal to zero, and its conditional variance is equal to $k_{jj}(1 - k_{jj})$.

- Autocovariance function of $h(\theta^i)$:

$$\begin{aligned} & \text{COV}(h(\theta^i), h(\theta^{(i-l)})) \\ &= (h(\tau_2) - h(\tau_1))^2 \pi_1(1 - \pi_1) \left(q - (1 - q) \frac{\pi_1}{1 - \pi_1} \right)^l \\ &= \mathbb{V}_\pi[h] \left(q - (1 - q) \frac{\pi_1}{1 - \pi_1} \right)^l \end{aligned}$$

- If $q = \pi_1$ then the autocovariances are equal to zero and the draws $h(\theta^i)$ are serially uncorrelated (in fact, in our simple discrete setting they are also independent).

- Define the Monte Carlo estimate

$$\bar{h}_N = \frac{1}{N} \sum_{i=1}^N h(\theta^i).$$

- Deduce from CLT

$$\sqrt{N}(\bar{h}_N - \mathbb{E}_\pi[h]) \implies N(0, \Omega(h)),$$


where $\Omega(h)$ is the long-run covariance matrix

$$\Omega(h) = \lim_{L \rightarrow \infty} \mathbb{V}_\pi[h] \left(1 + 2 \sum_{l=1}^L \frac{L-l}{L} \left(q - (1-q) \frac{\pi_1}{1-\pi_1} \right)^l \right).$$


- In turn, the asymptotic inefficiency factor is given by

$$\text{InEff}_\infty = \frac{\Omega(h)}{\mathbb{V}_\pi[h]} = 1 + 2 \lim_{L \rightarrow \infty} \sum_{l=1}^L \frac{L-l}{L} \left(q - (1-q) \frac{\pi_1}{1-\pi_1} \right)^l.$$


Autocorrelation Function of θ^i , $\pi_1 = 0.2$



Asymptotic Inefficiency $\text{InEff}_\infty, \pi_1 = 0.2$



Small Sample Variance $\mathbb{V}[\bar{h}_N]$ across Chains versus HAC Estimates of $\Omega(h)$



Solid line is 45-degree line.

- We discussed how to solve a DSGE model;
- and how to compute the likelihood function $p(Y|\theta)$ for a DSGE model.
- According to Bayes Theorem

$$p(\theta|Y) = \frac{p(Y|\theta)p(\theta)}{\int p(Y|\theta)p(\theta)d\theta}$$

- We want to generate draws from posterior...

Benchmark Random-Walk Metropolis-Hastings (RWMH) Algorithm for DSGE Models

- Initialization:
 - ① Use a numerical optimization routine to maximize the log posterior, which up to a constant is given by $\ln p(Y|\theta) + \ln p(\theta)$. Denote the posterior mode by $\hat{\theta}$.
 - ② Let $\hat{\Sigma}$ be the inverse of the (negative) Hessian computed at the posterior mode $\hat{\theta}$, which can be computed numerically.
 - ③ Draw θ^0 from $N(\hat{\theta}, c_0^2 \hat{\Sigma})$ or directly specify a starting value.
- Main Algorithm – For $i = 1, \dots, N$:
 - ① Draw ϑ from the proposal distribution $N(\theta^{i-1}, c^2 \hat{\Sigma})$.
 - ② Set $\theta^i = \vartheta$ with probability


$$\alpha(\vartheta|\theta^{i-1}) = \min \left\{ 1, \frac{p(Y|\vartheta)p(\vartheta)}{p(Y|\theta^{i-1})p(\theta^{i-1})} \right\}$$

and $\theta^i = \theta^{i-1}$ otherwise.

Benchmark RWMH Algorithm for DSGE Models


- Initialization steps can be modified as needed for particular application.
- If numerical optimization does not work well, one could let $\hat{\Sigma}$ be a diagonal matrix with prior variances on the diagonal.
- Or, $\hat{\Sigma}$ could be based on a preliminary run of a posterior sampler.
- It is good practice to run multiple chains based on different starting values.
- For the subsequent illustrations we chose $\hat{\Sigma} = \mathbb{V}_{\pi}[h]$, where the posterior variance matrix is obtained from a long MCMC run.

Observables for Small-Scale New Keynesian Model



Notes: Output growth per capita is measured in quarter-on-quarter (Q-o-Q) percentages. Inflation is CPI inflation in annualized Q-o-Q percentages. Federal funds rate is the average annualized effective funds rate for each quarter.

Convergence of Monte Carlo Average $\bar{\tau}_{N|N_0}$




Notes: The x -axis indicates the number of draws N . N_0 is set to 0, 25,000 and 50,000, respectively.

Posterior Estimates of DSGE Model Parameters

Parameter	Mean	[0.05, 0.95]	Parameter	Mean	[0.05,0.95]
τ	2.83	[1.95, 3.82]	ρ_r	0.77	[0.71, 0.82]
κ	0.78	[0.51, 0.98]	ρ_g	0.98	[0.96, 1.00]
ψ_1	1.80	[1.43, 2.20]	ρ_z	0.88	[0.84, 0.92]
ψ_2	0.63	[0.23, 1.21]	σ_r	0.22	[0.18, 0.26]
$r^{(A)}$	0.42	[0.04, 0.95]	σ_g	0.71	[0.61, 0.84]
$\pi^{(A)}$	3.30	[2.78, 3.80]	σ_z	0.31	[0.26, 0.36]
$\gamma^{(Q)}$	0.52	[0.28, 0.74]			

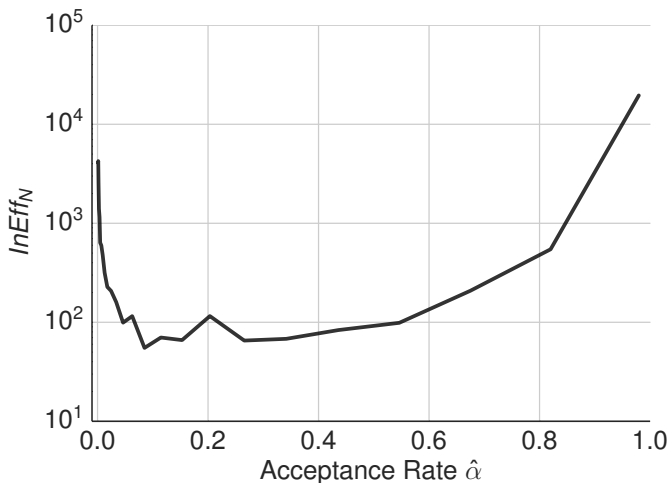
Notes: We generated $N = 100,000$ draws from the posterior and discarded the first 50,000 draws. Based on the remaining draws we approximated the posterior mean and the 5th and 95th percentiles.

Effect of Scaling Constant c




Notes: Results are based on $N_{run} = 50$ independent Markov chains. The acceptance rate (average across multiple chains), HAC-based estimate of $\ln \text{Eff}_\infty[\bar{\tau}]$ (average across multiple chains), and $\ln \text{Eff}_N[\bar{\tau}]$ are shown as a function of the scaling constant c .

Acceptance Rate $\hat{\alpha}$ versus Inaccuracy $\ln \text{Eff}_N$




Notes: $\ln \text{Eff}_N[\bar{\tau}]$ versus the acceptance rate $\hat{\alpha}$.

Impulse Responses of Exogenous Processes



Notes: The figure depicts pointwise posterior means and 90% credible bands. The responses are in percent relative to the initial level.

Parameter Transformations: Impulse Responses



Notes: The figure depicts pointwise posterior means and 90% credible bands. The responses of output are in percent relative to the initial level, whereas the responses of inflation and interest rates are in annualized percentages.

Challenges Due to Irregular Posteriors

- A stylized state-space model:

$$y_t = [1 \ 1]s_t, \quad s_t = \begin{bmatrix} \phi_1 & 0 \\ \phi_3 & \phi_2 \end{bmatrix} s_{t-1} + \begin{bmatrix} 1 \\ 0 \end{bmatrix} \epsilon_t, \quad \epsilon_t \sim iidN(0, 1).$$

where

- Structural parameters $\theta = [\theta_1, \theta_2]'$, domain is unit square.
- Reduced-form parameters $\phi = [\phi_1, \phi_2, \phi_3]'$

$$\phi_1 = \theta_1^2, \quad \phi_2 = (1 - \theta_1^2), \quad \phi_3 - \phi_2 = -\theta_1\theta_2.$$

Challenges Due to Irregular Posteriors

- $s_{1,t}$ looks like an exogenous technology process.
- $s_{2,t}$ evolves like an endogenous state variable, e.g., the capital stock.
- θ_2 is not identifiable if $\theta_1 = 0$ because θ_2 enters the model only multiplicatively.
- Law of motion of y_t is restricted ARMA(2,1) process:


$$(1 - \theta_1^2 L)(1 - (1 - \theta_1^2)L)y_t = (1 - \theta_1\theta_2 L)\epsilon_t.$$

- Given θ_1 and θ_2 , we obtain an observationally equivalent process by switching the values of the two roots of the autoregressive lag polynomial.
- Choose $\tilde{\theta}_1$ and $\tilde{\theta}_2$ such that


$$\tilde{\theta}_1 = \sqrt{1 - \theta_1^2}, \quad \tilde{\theta}_2 = \theta_1\theta_2/\tilde{\theta}_1.$$

Posteriors for Stylized State-Space Model

Local Identification Problem



Global Identification Problem



Notes: Intersections of the solid lines indicate parameter values that were used to generate the data from which the posteriors are constructed. Left panel: $\theta_1 = 0.1$ and $\theta_2 = 0.5$. Right panel: $\theta_1 = 0.8$, $\theta_2 = 0.3$.

Improvements to MCMC: Blocking

- In high-dimensional parameter spaces the RWMH algorithm generates highly persistent Markov chains.
- What's bad about persistence?

$$\sqrt{N}(\bar{h}_N - \mathbb{E}[\bar{h}_N]) \\ \implies N \left(0, \frac{1}{N} \sum_{i=1}^n \mathbb{V}[h(\theta^i)] + \frac{1}{N} \sum_{i=1}^N \sum_{j \neq i} \text{COV}[h(\theta^i), h(\theta^j)] \right).$$

- Potential Remedy:
 - Partition $\theta = [\theta_1, \dots, \theta_K]$.
 - Iterate over conditional posteriors $p(\theta_k | Y, \theta_{<-k})$.
- To reduce persistence of the chain, try to find partitions such that parameters are strongly correlated within blocks and weakly correlated across blocks or use random blocking.

Draw $\theta^0 \in \Theta$ and then for $i = 1$ to N :

- 1 Create a partition B^i of the parameter vector into N_{blocks} blocks $\theta_1, \dots, \theta_{N_{blocks}}$ via some rule (perhaps probabilistic), unrelated to the current state of the Markov chain.
- 2 For $b = 1, \dots, N_{blocks}$:

- 1 Draw $\vartheta_b \sim q(\cdot | [\theta_{<b}^i, \theta_b^{i-1}, \theta_{\geq b}^{i-1}])$.

- 2 With probability,

$$\alpha = \max \left\{ \frac{p([\theta_{<b}^i, \vartheta_b, \theta_{>b}^{i-1}] | Y) q(\theta_b^{i-1} | \theta_{<b}^i, \vartheta_b, \theta_{>b}^{i-1})}{p(\theta_{<b}^i, \theta_b^{i-1}, \theta_{>b}^{i-1} | Y) q(\vartheta_b | \theta_{<b}^i, \theta_b^{i-1}, \theta_{>b}^{i-1})}, 1 \right\},$$

set $\theta_b^i = \vartheta_b$, otherwise set $\theta_b^i = \theta_b^{i-1}$.

Random-Block MH Algorithm

- 1 Generate a sequence of random partitions $\{B^i\}_{i=1}^N$ of the parameter vector θ into N_{blocks} equally sized blocks, denoted by θ_b , $b = 1, \dots, N_{blocks}$ as follows:
 - 1 assign an $iidU[0, 1]$ draw to each element of θ ;
 - 2 sort the parameters according to the assigned random number;
 - 3 let the b 'th block consist of parameters $(b - 1)N_{blocks}, \dots, bN_{blocks}$.¹
- 2 Execute Algorithm Block MH Algorithm.

¹If the number of parameters is not divisible by N_{blocks} , then the size of a subset of the blocks has to be adjusted.

Metropolis-Adjusted Langevin Algorithm

- The proposal distribution of Metropolis-Adjusted Langevin (MAL) algorithm is given by

$$\begin{aligned}\mu(\theta^{i-1}) &= \theta^{i-1} + \frac{c_1}{2} M_1 \frac{\partial}{\partial \theta} \ln p(\theta^{i-1} | Y) \Big|_{\theta=\theta^{i-1}}, \\ \Sigma(\theta^{i-1}) &= c_2^2 M_2.\end{aligned}$$

that is θ^{i-1} is adjusted by a step in the direction of the gradient of the log posterior density function.

- One standard practice is to set $M_1 = M_2 = M$, with

$$M = - \left[\frac{\partial}{\partial \theta \partial \theta'} \ln p(\theta | Y) \Big|_{\theta=\hat{\theta}} \right]^{-1},$$

where $\hat{\theta}$ is the mode of the posterior distribution obtained using a numerical optimization routine.

Newton MH Algorithm

- Newton MH Algorithm replaces the Hessian evaluated at the posterior mode $\hat{\theta}$ by the Hessian evaluated at θ^{i-1} .
- The proposal distribution is given by

$$\begin{aligned}\mu(\theta^{i-1}) &= \theta^{i-1} - s \left[\frac{\partial}{\partial \theta \partial \theta'} \ln p(\theta | Y) \Big|_{\theta = \theta^{i-1}} \right]^{-1} \\ &\quad \times \frac{\partial}{\partial \theta} \ln p(\theta^{i-1} | Y) \Big|_{\theta = \theta^{i-1}} \\ \hat{\Sigma}(\theta^{i-1}) &= -c_2^2 \left[\frac{\partial}{\partial \theta \partial \theta'} \ln p(\theta | Y) \Big|_{\theta = \theta^{i-1}} \right]^{-1}.\end{aligned}$$

- It is useful to let s be independently of θ^{i-1} :

$$c_1 = 2s, \quad s \sim iidU[0, \bar{s}],$$


where \bar{s} is a tuning parameter.

Run Times and Tuning Constants for MH Algorithms

Algorithm	Run Time [hh:mm:ss]	Acceptance Rate	Tuning Constants
1-Block RWMH-I	00:01:13	0.28	$c = 0.015$
1-Block RWMH-V	00:01:13	0.37	$c = 0.400$
3-Block RWMH-I	00:03:38	0.40	$c = 0.070$
3-Block RWMH-V	00:03:36	0.43	$c = 1.200$
3-Block MAL	00:54:12	0.43	$c_1 = 0.400, c_2 = 0.750$
3-Block Newton MH	03:01:40	0.53	$\bar{s} = 0.700, c_2 = 0.600$


Notes: In each run we generate $N = 100,000$ draws. We report the fastest run time and the average acceptance rate across $N_{run} = 50$ independent Markov chains.

Autocorrelation Function of τ^i



Notes: The autocorrelation functions are computed based on a single run of each algorithm.

Inefficiency Factor $\text{InEff}_N[\bar{\tau}]$



Notes: The small-sample inefficiency factors are computed based on $N_{run} = 50$ independent runs of each algorithm.


IID Equivalent Draws Per Second

$$iid\text{-equivalent draws per second} = \frac{N}{\text{Run Time [seconds]}} \cdot \frac{1}{\text{InEff}_N}$$


Algorithm	Draws Per Second
1-Block RWMH-V	7.76
3-Block RWMH-V	5.65
3-Block MAL	1.24
3-Block RWMH-I	0.14
3-Block Newton MH	0.13
1-Block RWMH-I	0.04

Performance of Different MH Algorithms


RWMH-V (1 Block)




RWMH-V (3 Blocks)



MAL



Newton



Recall: Posterior Odds and Marginal Data Densities

- Posterior model probabilities can be computed as follows:

$$\pi_{i,T} = \frac{\pi_{i,0} p(Y|\mathcal{M}_i)}{\sum_j \pi_{j,0} p(Y|\mathcal{M}_j)}, \quad j = 1, \dots, 2, \quad (1)$$

- where

$$p(Y|\mathcal{M}) = \int p(Y|\theta, \mathcal{M}) p(\theta|\mathcal{M}) d\theta \quad (2)$$

- Note:

$$\ln p(Y_{1:T}|\mathcal{M}) = \sum_{t=1}^T \ln \int p(y_t|\theta, Y_{1:t-1}, \mathcal{M}) p(\theta|Y_{1:t-1}, \mathcal{M}) d\theta$$

- Posterior odds and Bayes Factor

$$\frac{\pi_{1,T}}{\pi_{2,T}} = \underbrace{\frac{\pi_{1,0}}{\pi_{2,0}}}_{\text{Prior Odds}} \times \underbrace{\frac{p(Y|\mathcal{M}_1)}{p(Y|\mathcal{M}_2)}}_{\text{Bayes Factor}} \quad (3)$$

Computation of Marginal Data Densities

- Reciprocal importance sampling:
 - Geweke's modified harmonic mean estimator
 - Sims, Waggoner, and Zha's estimator
- Chib and Jeliazkov's estimator
- For a survey, see Ardia, Hoogerheide, and van Dijk (2009).

- Reciprocal importance samplers are based on the following identity:

$$\frac{1}{p(Y)} = \int \frac{f(\theta)}{p(Y|\theta)p(\theta)} p(\theta|Y) d\theta, \quad (4)$$

where $\int f(\theta) d\theta = 1$.

- Conditional on the choice of $f(\theta)$ an obvious estimator is

$$\hat{p}_G(Y) = \left[\frac{1}{N} \sum_{i=1}^N \frac{f(\theta^i)}{p(Y|\theta^i)p(\theta^i)} \right]^{-1}, \quad (5)$$

where θ^i is drawn from the posterior $p(\theta|Y)$.

- Geweke (1999):

$$\begin{aligned} f(\theta) &= \tau^{-1} (2\pi)^{-d/2} |V_\theta|^{-1/2} \exp \left[-0.5(\theta - \bar{\theta})' V_\theta^{-1} (\theta - \bar{\theta}) \right] \\ &\times \left\{ (\theta - \bar{\theta})' V_\theta^{-1} (\theta - \bar{\theta}) \leq F_{\chi_d^2}^{-1}(\tau) \right\}. \end{aligned} \quad (6)$$

- Rewrite Bayes Theorem:

$$p(Y) = \frac{p(Y|\theta)p(\theta)}{p(\theta|Y)}. \quad (7)$$

- Thus,

$$\hat{p}_{CS}(Y) = \frac{p(Y|\tilde{\theta})p(\tilde{\theta})}{\hat{p}(\tilde{\theta}|Y)}, \quad (8)$$

where we replaced the generic θ in (7) by the posterior mode $\tilde{\theta}$.

- Use output of Metropolis-Hastings Algorithm.
- Proposal density for transition $\theta \mapsto \tilde{\theta}$: $q(\theta, \tilde{\theta} | Y)$.
- Probability of accepting proposed draw:

$$\alpha(\theta, \tilde{\theta} | Y) = \min \left\{ 1, \frac{p(\tilde{\theta} | Y)/q(\theta, \tilde{\theta} | Y)}{p(\theta | Y)/q(\tilde{\theta}, \theta | Y)} \right\}.$$

- Note that

$$\begin{aligned} & \int \alpha(\theta, \tilde{\theta} | Y) q(\theta, \tilde{\theta} | Y) p(\theta | Y) d\theta \\ &= \int \min \left\{ 1, \frac{p(\tilde{\theta} | Y)/q(\theta, \tilde{\theta} | Y)}{p(\theta | Y)/q(\tilde{\theta}, \theta | Y)} \right\} q(\theta, \tilde{\theta} | Y) p(\theta | Y) d\theta \\ &= p(\tilde{\theta} | Y) \int \min \left\{ \frac{p(\theta | Y)/q(\tilde{\theta}, \theta | Y)}{p(\tilde{\theta} | Y)/q(\theta, \tilde{\theta} | Y)}, 1 \right\} q(\tilde{\theta}, \theta | Y) d\theta \\ &= p(\tilde{\theta} | Y) \int \alpha(\tilde{\theta}, \theta | Y) q(\tilde{\theta}, \theta | Y) d\theta \end{aligned}$$

- Posterior density at the mode can be approximated as follows

$$\hat{p}(\tilde{\theta}|Y) = \frac{\frac{1}{N} \sum_{i=1}^N \alpha(\theta^i, \tilde{\theta}|Y) q(\theta^i, \tilde{\theta}|Y)}{\frac{1}{J} \sum_{j=1}^J \alpha(\tilde{\theta}, \theta^j|Y)}, \quad (9)$$

- $\{\theta^i\}$ are posterior draws obtained with the the M-H Algorithm;
- $\{\theta^j\}$ are additional draws from $q(\tilde{\theta}, \theta|Y)$ given the fixed value $\tilde{\theta}$.

MH-Based Marginal Data Density Estimates

Model	Mean($\ln \hat{p}(Y)$)	Std. Dev.($\ln \hat{p}(Y)$)
Geweke ($\tau = 0.5$)	-346.17	0.03
Geweke ($\tau = 0.9$)	-346.10	0.04
SWZ ($q = 0.5$)	-346.29	0.03
SWZ ($q = 0.9$)	-346.31	0.02
Chib and Jeliazkov	-346.20	0.40

Notes: Table shows mean and standard deviation of log marginal data density estimators, computed over $N_{run} = 50$ runs of the RWMH-V sampler using $N = 100,000$ draws, discarding a burn-in sample of $N_0 = 50,000$ draws. The SWZ estimator uses $J = 100,000$ draws to compute $\hat{\tau}$, while the CJ estimators uses $J = 100,000$ to compute the denominator of $\hat{p}(\tilde{\theta} | Y)$.