# Bayesian Inference for DSGE Models 

Lawrence J. Christiano

## Outline

- State space-observer form.
- convenient for model estimation and many other things.
- Bayesian inference
- Bayes' rule.
- Monte Carlo integation.
- MCMC algorithm.
- Laplace approximation


## State Space/Observer Form

- Compact summary of the model, and of the mapping between the model and data used in the analysis.
- Typically, data are available in log form. So, the following is useful:
- If $x$ is steady state of $x_{t}$ :

$$
\begin{aligned}
\hat{x}_{t} & \equiv \frac{x_{t}-x}{x} \\
& \Longrightarrow \frac{x_{t}}{x}=1+\hat{x}_{t} \\
& \Longrightarrow \log \left(\frac{x_{t}}{x}\right)=\log \left(1+\hat{x}_{t}\right) \approx \hat{x}_{t}
\end{aligned}
$$

- Suppose we have a model solution in hand: ${ }^{1}$

$$
\begin{aligned}
z_{t} & =A z_{t-1}+B s_{t} \\
s_{t} & =P s_{t-1}+\epsilon_{t}, E \epsilon_{t} \epsilon_{t}^{\prime}=D
\end{aligned}
$$

${ }^{1}$ Notation taken from solution lecture notes, http://faculty.wcas.northwestern.edu/~Ichrist/course/ Korea 2012 /lecture on solving rev.pdf

## State Space/Observer Form

- Suppose we have a model in which the date $t$ endogenous variables are capital, $K_{t+1}$, and labor, $N_{t}$ :

$$
z_{t}=\binom{\hat{K}_{t+1}}{\hat{N}_{t}}, s_{t}=\hat{\varepsilon}_{t}, \epsilon_{t}=e_{t} .
$$

- Data may include variables in $z_{t}$ and/or other variables.
- for example, suppose available data are $N_{t}$ and GDP, $y_{t}$ and production function in model is:

$$
y_{t}=\varepsilon_{t} K_{t}^{\alpha} N_{t}^{1-\alpha},
$$

so that

$$
\begin{aligned}
\hat{y}_{t} & =\hat{\varepsilon}_{t}+\alpha \hat{K}_{t}+(1-\alpha) \hat{N}_{t} \\
& =\left(\begin{array}{ll}
0 & 1-\alpha
\end{array}\right) z_{t}+\left(\begin{array}{ll}
\alpha & 0
\end{array}\right) z_{t-1}+s_{t}
\end{aligned}
$$

- From the properties of $\hat{y}_{t}$ and $\hat{N}_{t}$ :

$$
Y_{t}^{\text {data }}=\binom{\log y_{t}}{\log N_{t}}=\binom{\log y}{\log N}+\binom{\hat{y}_{t}}{\hat{N}_{t}}
$$

## State Space/Observer Form

- Model prediction for data:

$$
\begin{gathered}
Y_{t}^{\text {data }}=\binom{\log y}{\log N}+\binom{\hat{y}_{t}}{\hat{N}_{t}} \\
=\binom{\log y}{\log N}+\left[\begin{array}{cc}
0 & 1-\alpha \\
0 & 1
\end{array}\right] z_{t}+\left[\begin{array}{ll}
\alpha & 0 \\
0 & 0
\end{array}\right] z_{t-1}+\left[\begin{array}{l}
1 \\
0
\end{array}\right] s_{t} \\
=a+H \xi_{t} \\
\xi_{t}=\left(\begin{array}{c}
z_{t} \\
z_{t-1} \\
\hat{\varepsilon}_{t}
\end{array}\right), a=\left[\begin{array}{c}
\log y \\
\log N
\end{array}\right], H=\left[\begin{array}{ccccc}
0 & 1-\alpha & \alpha & 0 & 1 \\
0 & 1 & 0 & 0 & 0
\end{array}\right]
\end{gathered}
$$

- The Observer Equation may include measurement error, $w_{t}$ :

$$
Y_{t}^{\text {data }}=a+H \xi_{t}+w_{t}, E w_{t} w_{t}^{\prime}=R
$$

- Semantics: $\xi_{t}$ is the state of the system (do not confuse with the economic state $\left(K_{t}, \varepsilon_{t}\right)!$ ).


## State Space/Observer Form

- Law of motion of the state, $\xi_{t}$ (state-space equation):

$$
\begin{gathered}
\xi_{t}=F \xi_{t-1}+u_{t}, E u_{t} u_{t}^{\prime}=Q \\
\left(\begin{array}{c}
z_{t+1} \\
z_{t} \\
s_{t+1}
\end{array}\right)=\left[\begin{array}{ccc}
A & 0 & B P \\
I & 0 & 0 \\
0 & 0 & P
\end{array}\right]\left(\begin{array}{c}
z_{t} \\
z_{t-1} \\
s_{t}
\end{array}\right)+\left(\begin{array}{c}
B \\
0 \\
I
\end{array}\right) \epsilon_{t+1}, \\
u_{t}=\left(\begin{array}{c}
B \\
0 \\
I
\end{array}\right) \epsilon_{t,} Q=\left[\begin{array}{ccc}
B D B^{\prime} & 0 & B D \\
0 & 0 & 0 \\
D B^{\prime} & D
\end{array}\right], F=\left[\begin{array}{ccc}
A & 0 & B P \\
I & 0 & 0 \\
0 & 0 & P
\end{array}\right] .
\end{gathered}
$$

## State Space/Observer Form

$$
\xi_{t}=F \xi_{t-1}+u_{t}, E u_{t} u_{t}^{\prime}=Q,
$$

$$
Y_{t}^{\text {data }}=a+H \zeta_{t}+w_{t}, E w_{t} w_{t}^{\prime}=R
$$

- Can be constructed from model parameters

$$
\theta=(\beta, \delta, \ldots)
$$

SO

$$
F=F(\theta), Q=Q(\theta), a=a(\theta), H=H(\theta), R=R(\theta) .
$$

## Uses of State Space/Observer Form

- Estimation of $\theta$ and forecasting $\xi_{t}$ and $Y_{t}^{\text {data }}$
- Can take into account situations in which data represent a mixture of quarterly, monthly, daily observations.
- 'Data Rich' estimation. Could include several data measures (e.g., employment based on surveys of establishments and surveys of households) on a single model concept.
- Useful for solving the following forecasting problems:
- Filtering (mainly of technical interest in computing likelihood function):

$$
P\left[\xi_{t} \mid Y_{t-1}^{d a t a}, Y_{t-2}^{d d a t a}, \ldots, Y_{1}^{d a t a}\right], t=1,2, \ldots, T .
$$

- Smoothing:

$$
P\left[\xi_{t} \mid Y_{T}^{\text {data }}, \ldots, Y_{1}^{\text {data }}\right], t=1,2, \ldots, T
$$

- Example: 'real rate of interest' and 'output gap' can be recovered from $\xi_{t}$ using simple New Keynesian model.
- Useful for deriving a model's implications vector autoregressions


## Quick Review of Probability Theory

- Two random variables, $x \in\left(x_{1}, x_{2}\right)$ and $y \in\left(y_{1}, y_{2}\right)$.
- Joint distribution: $p(x, y)$

$$
=\begin{array}{ll}
y_{1} & 0.05 \\
y_{2} & 0.35 \\
\hline
\end{array}
$$

where

$$
p_{i j}=\operatorname{probability}\left(x=x_{i}, y=y_{j}\right) .
$$

- Restriction:

$$
\int_{x, y} p(x, y) d x d y=1
$$

## Quick Review of Probability Theory

- Joint distribution: $p(x, y)$
- Marginal distribution of $x: p(x)$

Probabilities of various values of $x$ without reference to the value of $y$ :

$$
p(x)=\left\{\begin{array}{ll}
p_{11}+p_{21}=0.40 & x=x_{1} \\
p_{12}+p_{22}=0.60 & x=x_{2}
\end{array} .\right.
$$

or,

$$
p(x)=\int_{y} p(x, y) d y
$$

## Quick Review of Probability Theory

- Joint distribution: $p(x, y)$
- Conditional distribution of $x$ given $y: p(x \mid y)$
- Probability of $x$ given that the value of $y$ is known

$$
p\left(x \mid y_{1}\right)= \begin{cases}p\left(x_{1} \mid y_{1}\right) & \frac{p_{11}}{p_{11}+p_{12}}=\frac{p_{11}}{p\left(y_{1}\right)}=\frac{0.05}{0.45}=0.11 \\ p\left(x_{2} \mid y_{1}\right) & \frac{p_{12}}{p_{11}+p_{12}}=\frac{p_{12}}{p\left(y_{1}\right)}=\frac{0.40}{0.45}=0.89\end{cases}
$$

or,

$$
p(x \mid y)=\frac{p(x, y)}{p(y)}
$$

## Quick Review of Probability Theory

- Joint distribution: $p(x, y)$

|  | $x_{1}$ | $x_{2}$ |  |
| :---: | :---: | :---: | :---: |
| $y_{1}$ | 0.05 | 0.40 | $p\left(y_{1}\right)=0.45$ |
| $y_{2}$ | 0.35 | 0.20 | $p\left(y_{2}\right)=0.55$ |
|  | $p\left(x_{1}\right)=0.40$ | $p\left(x_{2}\right)=0.60$ |  |

- Mode
- Mode of joint distribution (in the example):

$$
\operatorname{argmax}_{x, y} p(x, y)=\left(x_{2}, y_{1}\right)
$$

- Mode of the marginal distribution:

$$
\operatorname{argmax}_{x} p(x)=x_{2}, \operatorname{argmax}_{y} p(y)=y_{2}
$$

- Note: mode of the marginal and of joint distribution conceptually different.


## Maximum Likelihood Estimation

- State space-observer system:

$$
\begin{aligned}
\xi_{t+1} & =F \xi_{t}+u_{t+1}, E u_{t} u_{t}^{\prime}=Q \\
Y_{t}^{\text {data }} & =a_{0}+H \xi_{t}+w_{t}, E w_{t} w_{t}^{\prime}=R
\end{aligned}
$$

- Reduced form parameters, $\left(F, Q, a_{0}, H, R\right)$, functions of $\theta$.
- Choose $\theta$ to maximize likelihood, $p\left(Y^{\text {data }} \mid \theta\right)$ :

$$
\begin{aligned}
p\left(Y^{\text {data }} \mid \theta\right)= & p\left(Y_{1}^{\text {data }}, \ldots, Y_{T}^{\text {data }} \mid \theta\right) \\
= & p\left(Y_{1}^{\text {data }} \mid \theta\right) \times p\left(Y_{2}^{\text {data }} \mid Y_{1}^{\text {data }}, \theta\right) \\
& \times \cdots \times p\left(Y_{t}^{\text {data }} \mid Y_{t-1}^{\text {data }} \cdots Y_{1}^{\text {data }}, \theta\right)
\end{aligned}
$$

- Kalman filter straightforward (see, e.g., Hamilton's textbook).


## 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, $Y^{\text {data }}$.
- General property of probabilities:

$$
p\left(Y^{\text {data }}, \theta\right)=\left\{\begin{array}{c}
p\left(Y^{\text {data }} \mid \theta\right) \times p(\theta) \\
p\left(\theta \mid Y^{\text {data }}\right) \times p\left(Y^{\text {data }}\right)
\end{array}\right.
$$

which implies Bayes' rule:

$$
p\left(\theta \mid Y^{\text {data }}\right)=\frac{p\left(Y^{\text {data }} \mid \theta\right) p(\theta)}{p\left(Y^{\text {data }}\right)}
$$

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

## Bayesian Inference

- Report features of the posterior distribution, $p\left(\theta \mid Y^{\text {data }}\right)$.
- The value of $\theta$ that maximizes $p\left(\theta \mid Y^{d a t a}\right)$, 'mode' of posterior distribution.
- Compare marginal prior, $p\left(\theta_{i}\right)$, with marginal posterior of individual elements of $\theta, g\left(\theta_{i} \mid Y^{\text {data }}\right)$ :

$$
g\left(\theta_{i} \mid Y^{d a t a}\right)=\int_{\theta_{j \neq i}} p\left(\theta \mid Y^{d a t a}\right) d \theta_{j \neq i} \text { (multiple integration!!) }
$$

- Probability intervals about the mode of $\theta$ ('Bayesian confidence intervals'), need $g\left(\theta_{i} \mid Y^{d a t a}\right)$.
- Marginal likelihood for assessing model 'fit':

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

## 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

$$
\int_{a}^{b} f(x) d x,-\infty \leq a<b \leq \infty
$$

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

$$
\int_{a}^{b} f(x) d x=\int_{a}^{b} \frac{f(x)}{g(x)} g(x) d x=E \frac{f(x)}{g(x)}
$$

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 $i=1, \ldots, M$.
- step 2: evaluate $f\left(x_{i}\right) / g\left(x_{i}\right)$ and compute:

$$
\mu_{M} \equiv \frac{1}{M} \sum_{i=1}^{M} \frac{f\left(x_{i}\right)}{g\left(x_{i}\right)} \rightarrow_{M \rightarrow \infty} E \frac{f(x)}{g(x)} .
$$

- Exercise.
- Consider an integral where you have an analytic solution available, e.g., $\int_{0}^{1} x^{2} d x$.
- 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 $f\left(x_{i}\right) / g\left(x_{i}\right)$ over $i$ implies:

$$
\begin{gathered}
\operatorname{var}\left(\frac{1}{M} \sum_{i=1}^{M} \frac{f\left(x_{i}\right)}{g\left(x_{i}\right)}\right)=\frac{v_{M}}{M} \\
v_{M} \equiv \operatorname{var}\left(\frac{f\left(x_{i}\right)}{g\left(x_{i}\right)}\right) \simeq \frac{1}{M} \sum_{i=1}^{M}\left[\frac{f\left(x_{i}\right)}{g\left(x_{i}\right)}-\mu_{M}\right]^{2} .
\end{gathered}
$$

- Central Limit Theorem
- Estimate of $\int_{a}^{b} f(x) d x$ is a realization from a Nomal distribution with mean estimated by $\mu_{M}$ and variance, $v_{M} / M$.
- With 95\% probability,

$$
\mu_{M}-1.96 \times \sqrt{\frac{v_{M}}{M}} \leq \int_{a}^{b} f(x) d x \leq \mu_{M}+1.96 \times \sqrt{\frac{v_{M}}{M}}
$$

- Pick $g$ to minimize variance in $f\left(x_{i}\right) / g\left(x_{i}\right)$ 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, $\theta^{(1)}, \theta^{(2)}, \ldots, \theta^{(M)}$, of values of the $N \times 1$ vector of model parameters in such a way that

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

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


## MCMC Algorithm: setting up

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

$$
G(\theta)=\log p\left(Y^{d a t a} \mid \theta\right)+\log p(\theta)
$$

- Compute posterior mode:

$$
\theta^{*}=\arg \max _{\theta} G(\theta)
$$

- Compute the positive definite matrix, $V$ :

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

- 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 $\theta^{(r)}, x$,

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

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

$$
\lambda=\frac{p\left(Y_{\text {data }} \mid x\right) p(x)}{p\left(Y^{\text {data }} \mid \theta^{(r-1)}\right) p\left(\theta^{(r-1)}\right)}
$$

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

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

## 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, $\theta^{(1)}, \theta^{(2)}, \ldots, \theta^{(M)}$, 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).

Why a Low Acceptance Rate is Desirable


## MCMC Algorithm: using the Results

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

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

## 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

$$
\begin{aligned}
& f\left(\theta^{(1)}\right), \ldots, f\left(\theta^{(M)}\right) \\
& \rightarrow \quad E f(\theta) \simeq \bar{f} \equiv \frac{1}{M} \sum_{i=1}^{M} f\left(\theta^{(i)}\right),
\end{aligned}
$$

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

## MCMC: Remaining Issues

- In addition to the first and second moments already discused, 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:

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

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

By the law of large numbers,

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

## MCMC Algorithm: the Marginal Likelihood

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

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

$$
\frac{h\left(\theta^{(j)}\right)}{p\left(Y^{\text {data }} \mid \theta^{(j)}\right) p\left(\theta^{(j)}\right)}
$$

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,

$$
\begin{aligned}
& G(\theta) \equiv \log p\left(Y^{\text {data }} \mid \theta\right) p(\theta) \\
& p\left(Y^{\text {data }} \mid \theta\right) \sim \text { likelihood of data } \\
& p(\theta) \sim \text { prior on parameters } \\
& \theta^{*}{ }^{\sim} \text { maximum of } G(\theta) \text { (i.e., mode) }
\end{aligned}
$$

## Laplace Approximation

Second order Taylor series expansion of $G(\theta) \equiv \log \left[p\left(Y^{\text {data }} \mid \theta\right) p(\theta)\right]$ about $\theta=\theta^{*}$ :
$G(\theta) \approx G\left(\theta^{*}\right)+G_{\theta}\left(\theta^{*}\right)\left(\theta-\theta^{*}\right)-\frac{1}{2}\left(\theta-\theta^{*}\right)^{\prime} G_{\theta \theta}\left(\theta^{*}\right)\left(\theta-\theta^{*}\right)$,
where

$$
G_{\theta \theta}\left(\theta^{*}\right)=-\left.\frac{\partial^{2} \log p\left(Y^{\text {data }} \mid \theta\right) p(\theta)}{\partial \theta \partial \theta^{\prime}}\right|_{\theta=\theta^{*}}
$$

Interior optimality of $\theta^{*}$ implies:

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

Then:

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

## Laplace Approximation to Posterior Distribution

Property of Normal distribution:
$\int_{\theta} \frac{1}{(2 \pi)^{\frac{N}{2}}}\left|G_{\theta \theta}\left(\theta^{*}\right)\right|^{\frac{1}{2}} \exp \left\{-\frac{1}{2}\left(\theta-\theta^{*}\right)^{\prime} G_{\theta \theta}\left(\theta^{*}\right)\left(\theta-\theta^{*}\right)\right\} d \theta=1$
Then,

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

## Laplace Approximation

- Conclude:

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

- Laplace approximation to posterior distribution:

$$
\begin{aligned}
\frac{p\left(Y^{\text {data }} \mid \theta\right) p(\theta)}{p\left(Y^{\text {data }}\right)} \simeq & \frac{1}{(2 \pi)^{\frac{N}{2}}}\left|G_{\theta \theta}\left(\theta^{*}\right)\right|^{\frac{1}{2}} \\
& \times \exp \left\{-\frac{1}{2}\left(\theta-\theta^{*}\right)^{\prime} G_{\theta \theta}\left(\theta^{*}\right)\left(\theta-\theta^{*}\right)\right\}
\end{aligned}
$$

- So, posterior of $\theta_{i}$ (i.e., $g\left(\theta_{i} \mid Y^{\text {data }}\right)$ ) is approximately

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


gure 16 Priors and posteriors of estimated parameters of the medium-sized DSGE model.

## Modified Harmonic Mean Estimator of Marginal Likelihood

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

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

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 $p\left(Y^{\text {data }} \mid \theta^{(j)}\right) p\left(\theta^{(j)}\right)$ 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 $p\left(Y^{\text {data }} \mid\right.$ Model 1$)$ and $p\left(Y^{\text {data }} \mid\right.$ Model 2$)$
- Suppose $p\left(Y^{\text {data }} \mid\right.$ Model 1$)>p\left(Y^{\text {data }} \mid\right.$ 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.

