Bayesian Methods for Empirical Macroeconomics

Gary Koop

September 14, 2012

Bayesian Theory

- Reading: Chapter 1 of textbook and Appendix B, section B.1.
- Begin with general concepts in Bayesian theory before getting to specific models.
- If you know these general concepts you will never get lost.
- What does econometrician do? i) Estimate parameters in a model (e.g. regression coefficients), ii) Compare different models (e.g. hypothesis testing), iii) Prediction.
- Bayesian econometrics does these based on a few simple rules of probability.

- Let A and B be two events, $p(B|A)$ is the conditional probability of $B|A$. "summarizes what is known about B given A"
- Bayesians use this rule with $A =$ something known or assumed (e.g. the Data), B is something unknown (e.g. coefficients in a model).
- Let y be data, y^* be unobserved data (i.e. to be forecast), M_i for $i = 1, .., m$ be set of models each of which depends on some parameters, θ^i.
- Learning about parameters in a model is based on the posterior density: $p(\theta^i|M_i, y)$
- Model comparison based on posterior model probability: $p(M_i|y)$
- Prediction based on the predictive density $p(y^*|y)$.

Bayes Theorem

- I expect you know basics of probability theory from previous studies, see Appendix B of my textbook if you do not.
- *Definition: Conditional Probability*
- The conditional probability of A given B, denoted by $\Pr(A|B)$, is the probability of event A occurring given event B has occurred.
- *Theorem: Rules of Conditional Probability including Bayes' Theorem*
- Let A and B denote two events, then
- $\Pr(A|B) = \frac{\Pr(A,B)}{\Pr(B)}$ and
- $\Pr(B|A) = \frac{\Pr(A,B)}{\Pr(A)}$.

- These two rules can be combined to yield *Bayes' Theorem*:

$$\Pr(B|A) = \frac{\Pr(A|B)\Pr(B)}{\Pr(A)}.$$

- *Note:* Above is expressed in terms of two events, A and B. However, can be interpreted as holding for random variables, A and B with probability density functions replacing the $\Pr()$s in previous formulae.

Learning About Parameters in a Given Model (Estimation)

- Assume a single model which depends on parameters θ
- Want to figure out properties of the posterior $p(\theta|y)$
- It is convenient to use Bayes' rule to write the posterior in a different way.
- Bayes' rule lies at the heart of Bayesian econometrics:

$$p(B|A) = \frac{p(A|B)p(B)}{p(A)}.$$

- Replace B by θ and A by y to obtain:

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{p(y)}.$$

- Bayesians treat $p(\theta|y)$ as being of fundamental interest: "Given the data, what do we know about θ?".
- Treatment of θ as a random variable is controversial among some econometricians.
- Competitor to Bayesian econometrics, called *frequentist econometrics*, says that θ is not a random variable.
- For estimation can ignore the term $p(y)$ since it does not involve θ:

$$p(\theta|y) \propto p(y|\theta)p(\theta).$$

- $p(\theta|y)$ is referred to as the *posterior density*
- $p(y|\theta)$ is the *likelihood function*
- $p(\theta)$ as the *prior density*.
- "posterior is proportional to likelihood times prior".

- $p(\theta)$, does not depend on the data. It contains any non-data information available about θ.
- Prior information is controversial aspect since it sounds unscientific.
- Bayesian answers (to be elaborated on later):
- i) Often we do have prior information and, if so, we should include it (more information is good)
- ii) Can work with "noninformative" priors
- iii) Can use "empirical Bayes" methods which estimate prior from the data
- iv) Training sample priors
- v) Bayesian estimators often have better frequentist properties than frequentist estimators (e.g. results due to Stein show MLE is inadmissible – but Bayes estimators are admissible)
- vi) Prior sensitivity analysis

Prediction in a Single Model

- Prediction based on the *predictive density* $p(y^*|y)$
- Since a marginal density can be obtained from a joint density through integration:

$$p(y^*|y) = \int p(y^*, \theta|y) d\theta.$$

- Term inside integral can be rewritten as:

$$p(y^*|y) = \int p(y^*|y, \theta) p(\theta|y) d\theta.$$

- Prediction involves the posterior and $p(y^*|y, \theta)$ (more description provided later)

Model Comparison (Hypothesis testing)

- Models denoted by M_i for $i = 1, .., m$. M_i depends on parameters θ^i.
- *Posterior model probability* is $p(M_i|y)$.
- Using Bayes rule with $B = M_i$ and $A = y$ we obtain:

$$p(M_i|y) = \frac{p(y|M_i)p(M_i)}{p(y)}$$

- $p(M_i)$ is referred to as the *prior model probability*.
- $p(y|M_i)$ is called the *marginal likelihood.*

- How is marginal likelihood calculated?
- Posterior can be written as:

$$p(\theta^i|y, M_i) = \frac{p(y|\theta^i, M_i)p(\theta^i|M_i)}{p(y|M_i)}$$

- Integrate both sides with respect to θ^i, use fact that $\int p(\theta^i|y, M_i)d\theta^i = 1$ and rearrange:

$$p(y|M_i) = \int p(y|\theta^i, M_i)p(\theta^i|M_i)d\theta^i.$$

- Note: marginal likelihood depends only on the prior and likelihood.

- *Posterior odds ratio* compares two models:

$$PO_{ij} = \frac{p(M_i|y)}{p(M_j|y)} = \frac{p(y|M_i)p(M_i)}{p(y|M_j)p(M_j)}.$$

- Note: $p(y)$ is common to both models, no need to calculate.
- Can use fact that $p(M_1|y) + p(M_2|y) + ... + p(M_m|y) = 1$ and PO_{ij} to calculate the posterior model probabilities.
- E.g. if $m = 2$ models:

$$p(M_1|y) + p(M_2|y) = 1$$

$$PO_{12} = \frac{p(M_1|y)}{p(M_2|y)}$$

- imply

$$p(M_1|y) = \frac{PO_{12}}{1 + PO_{12}}$$

$$p(M_2|y) = 1 - p(M_1|y).$$

- The *Bayes Factor* is:

$$BF_{ij} = \frac{p(y|M_i)}{p(y|M_j)}.$$

Summary

- These few pages have outlined all the basic theoretical concepts required for the Bayesian to learn about parameters, compare models and predict.
- This is an enormous advantage: Once you accept that unknown things (i.e. θ, M_i and y^*) are random variables, the rest of Bayesian approach is non-controversial.
- What are going to do in rest of this course?
- See how these concepts work in some models of interest.
- First the regression model (very briefly)
- Then time series models of interest for macroeconomics
- Bayesian computation.

Bayesian Computation

- How do you present results from a Bayesian empirical analysis?
- $p(\theta|y)$ is a p.d.f. Especially if θ is a vector of many parameters cannot present a graph of it.
- Want features analogous to frequentist point estimates and confidence intervals.
- A common point estimate is the mean of the posterior density (or *posterior mean*).
- Let θ be a vector with k elements, $\theta = (\theta_1, .., \theta_k)'$. The posterior mean of any element of θ is:

$$E(\theta_i|y) = \int \theta_i p(\theta|y) d\theta.$$

- Common measure of dispersion is the *posterior standard deviation* (square root of *posterior variance*)
- Posterior variance:

$$var(\theta_i|y) = E(\theta_i^2|y) - \{E(\theta_i|y)\}^2,$$

- This requires calculating another expected value:

$$E(\theta_i^2|y) = \int \theta_i^2 p(\theta|y) d\theta.$$

- Many other possible features of interest. E.g. what is probability that a coefficient is positive?

$$p(\theta_i \geq 0|y) = \int_0^\infty p(\theta_i|y) d\theta_i$$

- All of these posterior features have the form:

$$E\left[g\left(\theta\right)|y\right]=\int g(\theta)p(\theta|y)d\theta,$$

- where $g(\theta)$ is a *function of interest*.
- All these features have integrals in them. Marginal likelihood and predictive density also involved integrals.
- Apart from a few simple cases, it is not possible to evaluate these integrals analytically, and we must turn to the computer.

Posterior Simulation

- The integrals involved in Bayesian analysis are usually evaluated using simulation methods.
- Will use several methods later on. Here we provide some intuition.
- Frequentist asymptotic theory uses Laws of Large Numbers (LLN) and a Central Limit Theorems (CLT).
- A typical LLN: "consider a random sample, $Y_1, ..Y_N$, as N goes to infinity, the average converges to its expectation" (e.g. $\overline{Y} \rightarrow \mu$)
- Bayesians use LLN: "consider a random sample from the posterior, $\theta^{(1)}, ..\theta^{(S)}$, as S goes to infinity, the average of these converges to $E\left[\theta|y\right]$"
- Note: Bayesians use asymptotic theory, but asymptotic in S (under control of researcher) not N

- Example: Monte Carlo integration.
- Let $\theta^{(s)}$ for $s = 1, .., S$ be a random sample from $p(\theta|y)$ and define

$$\widehat{g}_S = \frac{1}{S} \sum_{s=1}^{S} g\left(\theta^{(s)}\right),$$

- then $\widehat{g}_S$ converges to $E\left[g\left(\theta\right)|y\right]$ as S goes to infinity.
- Monte Carlo integration approximates $E\left[g\left(\theta\right)|y\right]$, but only if S were infinite would the approximation error be zero.
- We can choose any value for S (but larger values of S will increase computational burden).
- To gauge size of approximation error, use a CLT to obtain numerical standard error.

- Most Bayesians write own programs (e.g. using Gauss, Matlab, R or C++) to do posterior simulation
- BUGS (Bayesian Analysis Using Gibbs Sampling) is a popular Bayesian package, but only has limited set of models (or require substantial programming to adapt to other models)
- Bayesian work cannot (easily) be done in standard econometric packages like Microfit, Eviews or Stata.
- I have a Matlab website for VARs, TVP-VARs and TVP-FAVARs (see my website)
- Peter Rossi has an R package for marketing and microeconomic applications
- http://faculty.chicagobooth.edu/peter.rossi/research/bsm.html
- Jim LeSage's Econometrics toolbox (Matlab)
- http://www.spatial-econometrics.com/
- Many more using R see http://cran.r-project.org/web/views/Bayesian.html

Bayesian Analysis of the Normal Linear Regression Model

- Regression model can be written as:

$$y = X\beta + \varepsilon.$$

- ε, y are $N \times 1$ vectors
- β is $k \times 1$ vector
- X is $N \times k$ matrix

The Likelihood Function

- Likelihood can be derived under the classical assumptions:
- ε is $N(0_N, h^{-1}I_N)$ where $h = \sigma^{-2}$.
- All elements of X are either fixed (i.e. not random variables).
- Exercise 10.1, Bayesian Econometric Methods shows that likelihood function can be written in terms of OLS quantities:

$$\nu = N - k,$$

$$\widehat{\beta} = \left(X'X\right)^{-1} X'y$$

$$s^2 = \frac{\left(y - X\widehat{\beta}\right)' \left(y - X\widehat{\beta}\right)}{\nu}$$

- Likelihood function:

$$p(y|\beta, h) = \frac{1}{(2\pi)^{\frac{N}{2}}}$$
$$\left\{h^{\frac{1}{2}} \exp\left[-\frac{h}{2}\left(\beta - \widehat{\beta}\right)' X'X \left(\beta - \widehat{\beta}\right)\right]\right\} \left\{h^{\frac{\nu}{2}} \exp\left[-\frac{h\nu}{2s^{-2}}\right]\right\}$$

The Prior

- Common starting point is natural conjugate Normal-Gamma prior
- β conditional on h is now multivariate Normal:

$$\beta|h \sim N(\underline{\beta}, h^{-1}\underline{V})$$

- Prior for error precision h is Gamma

$$h \sim G(\underline{s}^{-2}, \underline{\nu})$$

- $\underline{\beta}$, $\underline{V}$, $\underline{s}^{-2}$ and $\underline{\nu}$ a prior hyperparameter values chosen by the researcher
- Notation: Normal-Gamma distribution

$$\beta, h \sim NG\left(\underline{\beta}, \underline{V}, \underline{s}^{-2}, \underline{\nu}\right).$$

- Noninformative prior is limiting case, leads to OLS quantities

The Posterior

- Multiply likelihood by prior and collecting terms (see Bayesian Econometrics Methods Exercise 10.1).
- Posterior is

$$\beta, h|y \sim NG\left(\overline{\beta}, \overline{V}, \overline{s}^{-2}, \overline{\nu}\right)$$

- where

$$\overline{V} = \left(\underline{V}^{-1} + X'X\right)^{-1},$$

$$\overline{\beta} = \overline{V}\left(\underline{V}^{-1}\underline{\beta} + X'X\widehat{\beta}\right)$$

$$\overline{\nu} = \underline{\nu} + N$$

and $\overline{s}^{-2}$ is defined implicitly through

$$\overline{\nu}\overline{s}^2 = \underline{\nu}\underline{s}^2 + \nu s^2 + \left(\widehat{\beta} - \underline{\beta}\right)'\left[\underline{V} + \left(X'X\right)^{-1}\right]^{-1}\left(\widehat{\beta} - \underline{\beta}\right).$$

- Marginal posterior for β: multivariate t distribution:

$$\beta|y \sim t\left(\overline{\beta}, \overline{s}^2\overline{V}, \overline{\nu}\right),$$

- Useful results for estimation:

$$E(\beta|y) = \overline{\beta}$$

- $$var(\beta|y) = \frac{\overline{\nu}\overline{s}^2}{\overline{\nu}-2}\overline{V}.$$

- Intuition: Posterior mean and variance are weighted average of information in the prior and the data.

Model Comparison

- Case 1: M_1 imposes a linear restriction and M_2 does not (nested).
- Case 2: $M_1 : y = X_1\beta_{(1)} + \varepsilon_1$ and $M_2 : y = X_2\beta_{(2)} + \varepsilon_2$, where X_1 and X_2 contain different explanatory variables (non-nested).
- Both cases can be handled by defining models as (for $j = 1, 2$):

$$M_j : y_j = X_j\beta_{(j)} + \varepsilon_j$$

- Non-nested model comparison involves $y_1 = y_2$.
- Nested model comparison defines M_2 as unrestricted regression. M_1 imposes the restriction can involve a redefinition of explanatory and dependent variable.

Example: Nested Model Comparison

- M_2 is unrestricted model

$$y = \beta_1 + \beta_2 x_2 + \beta_3 x_3 + \varepsilon$$

- M_1 restricts $\beta_3 = 1$, can be written:

$$y - x_3 = \beta_1 + \beta_2 x_2 + \varepsilon$$

- M_1 has dependent variable $y - x_3$ and intercept and x_2 are explanatory variables

- Marginal likelihood is (for $j = 1, 2$):

$$p(y_j|M_j) = c_j \left(\frac{|\overline{V}_j|}{|\underline{V}_j|} \right)^{\frac{1}{2}} \left(\overline{\nu}_j \overline{s}_j^2 \right)^{-\frac{\overline{\nu}_j}{2}}$$

- c_j is constant depending on prior hyperparameters, etc.
-
$$PO_{12} = \frac{c_1 \left(\frac{|\overline{V}_1|}{|\underline{V}_1|} \right)^{\frac{1}{2}} \left(\overline{\nu}_1 \overline{s}_1^2 \right)^{-\frac{\overline{\nu}_1}{2}} p(M_1)}{c_2 \left(\frac{|\overline{V}_2|}{|\underline{V}_2|} \right)^{\frac{1}{2}} \left(\overline{\nu}_2 \overline{s}_2^2 \right)^{-\frac{\overline{\nu}_2}{2}} p(M_2)}$$

- Posterior odds ratio depends on the prior odds ratio and contains rewards for model fit, coherency between prior and data information and parsimony.

Model Comparison with Noninformative Priors

- Important rule: *When comparing models using posterior odds ratios, it is acceptable to use noninformative priors over parameters which are common to all models. However, informative, proper priors should be used over all other parameters.*
- If we set $\underline{\nu}_1 = \underline{\nu}_2 = 0$. Posterior odds ratio still has a sensible interpretation.
- Noninformative prior for h_1 and h_2 is fine (these parameters common to both models)
- But noninformative priors for $\beta_{(j)}$'s causes problems which occur largely when $k_1 \neq k_2$. (Exercise 10.4 of Bayesian Econometric Methods)
- E.g. noninformative prior for $\beta_{(j)}$ based on $\underline{V}_j^{-1} = cI_{k_j}$ and letting $c \rightarrow 0$. Since $|\underline{V}_j| = \frac{1}{c^{k_j}}$ terms involving k_j do not cancel out.
- If $k_1 < k_2$, PO_{12} becomes infinite, while if $k_1 > k_2$, PO_{12} goes to zero.

Prediction

- Want to predict:

$$y^* = X^*\beta + \varepsilon^*$$

- Remember, prediction is based on:

$$p(y^*|y) = \int\int p(y^*|y,\beta,h)\, p(\beta,h|y)\, d\beta dh.$$

- The resulting predictive:

$$y^*|y \sim t\left(X^*\overline{\beta}, \overline{s}^2\left\{I_T + X^*\overline{V}X^{*\prime}\right\}, \overline{\nu}\right)$$

- Model comparison, prediction and posterior inference about β can all be done analytically.
- So no need for posterior simulation in this model.
- However, let us illustrate Monte Carlo integration in this model.

Monte Carlo Integration

- Remember the basic LLN we used for Monte Carlo integration
- Let $\beta^{(s)}$ for $s = 1, .., S$ be a random sample from $p(\beta|y)$ and $g(.)$ be any function and define

$$\widehat{g}_S = \frac{1}{S}\sum_{r=1}^{S} g\left(\beta^{(s)}\right)$$

- then $\widehat{g}_S$ converges to $E\left[g(\beta)|y\right]$ as S goes to infinity.
- How would you write a computer program which did this?

- *Step 1:* Take a random draw, $\beta^{(s)}$ from the posterior for β using a random number generator for the multivariate t distribution.
- *Step 2:* Calculate $g\left(\beta^{(s)}\right)$ and keep this result.
- *Step 3:* Repeat Steps 1 and 2 S times.
- *Step 4:* Take the average of the S draws $g\left(\beta^{(1)}\right), ..., g\left(\beta^{(S)}\right)$.
- These steps will yield an estimate of $E\left[g(\beta)|y\right]$ for any function of interest.
- Remember: Monte Carlo integration yields only an approximation for $E\left[g(\beta)|y\right]$ (since you cannot set $S = \infty$).
- By choosing S, can control the degree of approximation error.
- Using a CLT we can obtain 95% confidence interval for $E[g(\beta)|y]$
- Or a numerical standard error can be reported.

Summary

- So far we have worked with Normal linear regression model using natural conjugate prior
- This meant posterior, marginal likelihood and predictive distributions had analytical forms
- But with other priors do not get analytical results.
- Next we try a new prior so as to introduce an important tool for posterior computation: the Gibbs sampler.
- The Gibbs sampler is a special type of Markov Chain Monte Carlo (MCMC) algorithm.

Normal Linear Regression Model with Independent Normal-Gamma Prior

- Keep the Normal linear regression model (under the classical assumptions) as before.
- Likelihood function presented above
- Parameters of model are β and h.

The Prior

- Before we had conjugate prior where $p(\beta|h)$ was Normal density and $p(h)$ Gamma density.
- Now use similar prior, but assume prior independence between β and h.
- $p(\beta, h) = p(\beta)\, p(h)$ with $p(\beta)$ being Normal and $p(h)$ being Gamma:

$$\beta \sim N\left(\underline{\beta}, \underline{V}\right)$$

and

$$h \sim G(\underline{s}^{-2}, \underline{\nu})$$

Key difference: now $\underline{V}$ is now the prior covariance matrix of β, with conjugate prior we had $var(\beta|h) = h^{-1}\underline{V}$.

The Posterior

- The posterior is proportional to prior times the likelihood.
- The joint posterior density for β and h does not take form of any well-known and understood density – cannot be directly used for posterior inference.
- However, conditional posterior for β (i.e. conditional on h) takes a simple form:

$$\beta | y, h \sim N\left(\overline{\beta}, \overline{V}\right)$$

- where

$$\overline{V} = \left(\underline{V}^{-1} + hX'X\right)^{-1}$$

$$\overline{\beta} = \overline{V}\left(\underline{V}^{-1}\underline{\beta} + hX'y\right)$$

- Conditional posterior for h takes simple form:

$$h|y,\beta \sim G(\overline{s}^{-2},\overline{\nu})$$

where

$$\overline{\nu} = N + \underline{\nu}$$

and

$$\overline{s}^2 = \frac{(y - X\beta)'(y - X\beta) + \underline{\nu}\underline{s}^2}{\overline{\nu}}$$

- Econometrician is interested in $p(\beta, h|y)$ (or $p(\beta|y)$), NOT the posterior conditionals, $p(\beta|y,h)$ and $p(h|y,\beta)$.
- Since $p(\beta, h|y) \neq p(\beta|y,h)\,p(h|y,\beta)$, the conditional posteriors do not directly tell us about $p(\beta, h|y)$.
- But, there is a posterior simulator, called the *Gibbs sampler*, which uses conditional posteriors to produce random draws, $\beta^{(s)}$ and $h^{(s)}$ for $s = 1, .., S$, which can be averaged to produce estimates of posterior properties just as with Monte Carlo integration.

The Gibbs Sampler

- Gibbs sampler is powerful tool for posterior simulation used in many econometric models.
- We will motivate general ideas before returning to regression model
- General notation: θ is a $p-$vector of parameters and $p(y|\theta)$, $p(\theta)$ and $p(\theta|y)$ are the likelihood, prior and posterior, respectively.
- Let θ be partitioned into *blocks* as $\theta = \left(\theta'_{(1)}, \theta'_{(2)}, .., \theta'_{(B)}\right)'$. E.g. in regression model set $B = 2$ with $\theta_{(1)} = \beta$ and $\theta_{(2)} = h$.

- Intuition: i) Monte Carlo integration takes draws from $p(\theta|y)$ and averages them to produce estimates of $E[g(\theta)|y]$ for any function of interest $g(\theta)$.
- ii) In many models, it is not easy to draw from $p(\theta|y)$. However, it often is easy to draw from $p\left(\theta_{(1)}|y,\theta_{(2)},..,\theta_{(B)}\right)$, $p\left(\theta_{(2)}|y,\theta_{(1)},\theta_{(3)}..,\theta_{(B)}\right)$, ..., $p\left(\theta_{(B)}|y,\theta_{(1)},..,\theta_{(B-1)}\right)$.
- Note: Preceding distributions are *full conditional posterior distributions* since they define a posterior for each block conditional on all other blocks.
- iii) Drawing from the full conditionals will yield a sequence $\theta^{(1)},\theta^{(2)},..,\theta^{(s)}$ which can be averaged to produce estimates of $E[g(\theta)|y]$ in the same manner as Monte Carlo integration.
- This is called Gibbs sampling

More motivation for the Gibbs sampler

- Regression model with $B = 2$: β and h
- Suppose that you have one random draw from $p(\beta|y)$. Call this draw $\beta^{(0)}$.
- Since $p(\beta, h|y) = p(h|y, \beta)\, p(\beta|y)$, a draw from $p\left(h|y, \beta^{(0)}\right)$ is a valid draw of h. Call this $h^{(1)}$.
- Since $p(\beta, h|y) = p(\beta|y, h)\, p(h|y)$, a random draw from $p\left(\beta|y, h^{(1)}\right)$ is a valid draw of β. Call this $\beta^{(1)}$
- Hence, $\left(\beta^{(1)}, h^{(1)}\right)$ is a valid draw from $p(\beta, h|y)$.
- You can continue this reasoning indefinitely producing $\left(\beta^{(s)}, h^{(s)}\right)$ for $s = 1, .., S$

- Hence, if you can successfully find $\beta^{(0)}$, then sequentially drawing $p(h|y,\beta)$ and $p(\beta|y,h)$ will give valid draws from posterior.
- Problem with above strategy is that it is not possible to find such an initial draw $\beta^{(0)}$.
- If we knew how to easily take random draws from $p(\beta|y)$, we could use this and $p(h|\beta,y)$ to do Monte Carlo integration and have no need for Gibbs sampling.
- However, it can be shown that subject to weak conditions, the initial draw $\beta^{(0)}$ does not matter: Gibbs sampler will converge to a sequence of draws from $p(\beta,h|y)$.
- In practice, choose $\beta^{(0)}$ in some manner and then run the Gibbs sampler for S replications.
- Discard S_0 initial draws ("the *burn-in*") and remaining S_1 used to estimate $E[g(\theta)|y]$

Why is Gibbs sampling so useful?

- In Normal linear regression model with independent Normal-Gamma prior Gibbs sampler is easy
- $p(\beta|y, h)$ is Normal and $p(h|y, \beta)$ and Gamma (easy to draw from)
- Huge number of other models have hard joint posterior, but easy posterior conditionals
- tobit, probit, stochastic frontier model, Markov switching model, threshold autoregressive, smooth transition threshold autoregressive, other regime switching models, state space models, some semiparametric regression models, etc etc etc.
- Also models of form I will now discuss

Summary

- This lecture shows how Bayesian ideas work in familiar context (regression model)
- Occasionally analytical results are available (no need for posterior simulation)
- Usually posterior simulation is required.
- Monte Carlo integration is simplest, but rarely possible to use it.
- Gibbs sampling (and related MCMC) methods can be used for estimation and prediction for a wide variety of models
- Note: There are methods for calculating marginal likelihoods using Gibbs sampler output
- Now we move on to models of interest for the empirical macroeconomist...