Consumption Function Case Study (David Giles)

- We want to estimate a simple aggregate consumption function.
- Annual U.S. data for the period 1950 to 1985.
- The model is: $C_t = \alpha + \beta Y_t + \varepsilon_t$; $\varepsilon_t \sim i.i.d.N[0, \sigma^2]$
- Take deviations about sample means:

$$c_t = \beta y_t + \varepsilon_t \quad ; \quad \varepsilon_t \sim i.i.d.N[0, \sigma^2]$$

- Prior information:

(i) $0 \leq \beta \leq 1$; (m.p.c.)

$E[\beta] = 0.8$; $V[\beta] = 0.001$

(ii) $0 < \sigma < \infty$

- Construct the Prior p.d.f. for the parameters:

$$p(\beta, \sigma) = p(\beta)p(\sigma) \quad ; \quad \textit{independent information}$$

$$p(\sigma) \propto 1/\sigma \quad ; \quad \textit{"diffuse", non-informative}$$

$$p(\beta) \propto \beta^{a-1}(1-\beta)^{b-1} \quad ; \quad a, b > 0 \; ; \; \textit{Beta density}$$

- $E[\beta] = a/(a+b) \; = 0.8$

  $V[\beta] = ab/[(a+b)^2(a+b+1)] = 0.001$

  $\Rightarrow a = 127.2 \; ; \; b = 31.8$

- So, $p(\beta, \sigma) \propto \sigma^{-1}\beta^{126.2}(1-\beta)^{30.8}$

- We don't need to worry about the proportionality constant - this will be dealt with once we have the Posterior p.d.f..

Prior p.d.f. for Beta

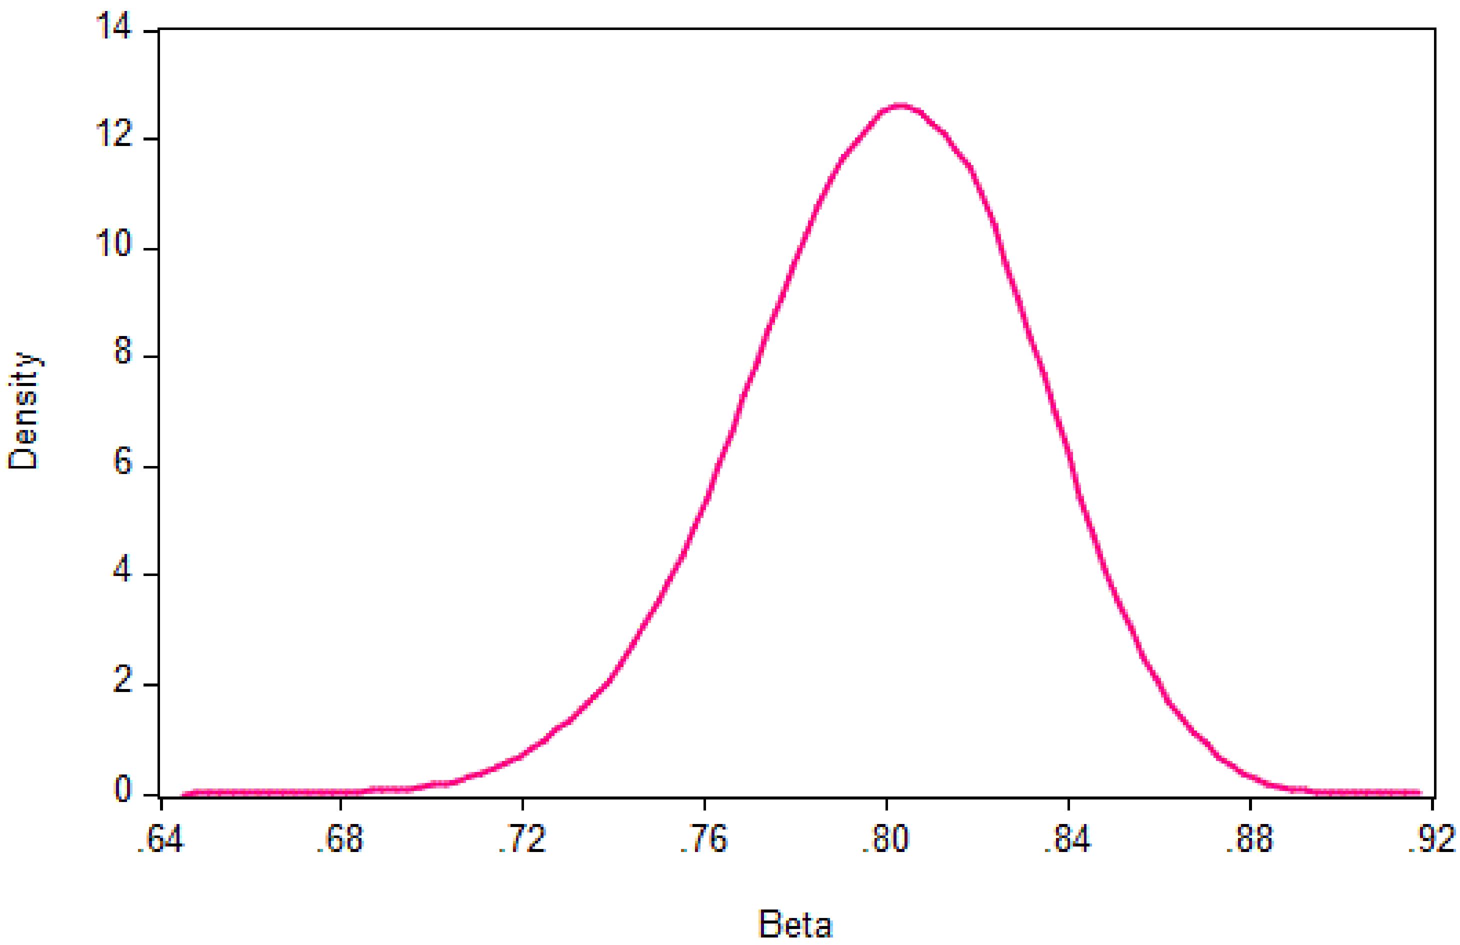

- Now we can look at the data and formulate the Likelihood Function

$$p(\boldsymbol{c}|\beta, \boldsymbol{\sigma}) = \left(\frac{1}{\boldsymbol{\sigma}\sqrt{\boldsymbol{2\pi}}}\right)^{\boldsymbol{n}} exp\left\{-\frac{1}{2\sigma^2}\sum_{t=1}^{n}(c_t - \beta y_t)^2\right\}$$

- Apply Bayes' Theorem:

$$p(\beta, \sigma|\boldsymbol{c}) \propto \sigma^{-1}\beta^{126.2}(1-\beta)^{30.8}\left(\frac{1}{\boldsymbol{\sigma}\sqrt{\boldsymbol{2\pi}}}\right)^{\boldsymbol{n}} exp\left\{-\frac{1}{2\sigma^2}\sum_{t=1}^{n}(c_t - \beta y_t)^2\right\}$$

$$\propto \sigma^{-(n+1)}\beta^{126.2}(1-\beta)^{30.8} exp\left\{-\frac{1}{2\sigma^2}\sum_{t=1}^{n}(c_t - \beta y_t)^2\right\}$$

- *This is not the kernel of any standard density function*.
- So, now we have to:

(i) Obtain the Marginal Posteriors for β and σ (& Normalize them).

(ii) Use these to get Bayes' estimators of β and σ.

- All of this requires some **Integration !**
- We can marginalize the Joint Posterior to get the Marginal Posterior for β analytically (*the details of the integration are provided later on*):

$$p(\beta|\boldsymbol{c}) = \int_0^\infty p(\beta, \sigma|\boldsymbol{c}) d\sigma$$

$$\propto \int_0^\infty \sigma^{-(n+1)} \beta^{126.2} (1-\beta)^{30.8} \, exp\left\{-\frac{1}{2\sigma^2} \sum_{t=1}^{n} (c_t - \beta y_t)^2\right\} d\sigma$$

$$\propto \beta^{126.2} (1-\beta)^{30.8} \left\{\sum_{t=1}^{n} (c_t - \beta y_t)^2\right\}^{-n/2}$$

- Now we need to *normalize* this Marginal Posterior p.d.f. for β.
- It should *integrate to one*, over the interval $0 < \beta < 1$.

- This means that we have to find the "normalizing constant",

$$k = 1/\left[\int_0^1 \beta^{126.2}(1-\beta)^{30.8}\{\sum_{t=1}^{n}(c_t - \beta y_t)^2\}^{-n/2}\, d\beta\right]$$

- This can be obtained quite easily by *numerical integration*, given the range of integration.

- Then,

$$p(\beta|\boldsymbol{c}) = k\beta^{126.2}(1-\beta)^{30.8}\left\{\sum_{t=1}^{n}(c_t - \beta y_t)^2\right\}^{-n/2}$$

- Recall our various Loss functions:

(i) $L_1(\theta, \hat{\theta}) = c_1(\theta - \hat{\theta})^2$; $c_1 > 0$

(ii) $L_2(\theta, \hat{\theta}) = c_2|\theta - \hat{\theta}|$; $c_2 > 0$

(iii) $L_3(\theta, \hat{\theta}) = c_3$; if $|\theta - \hat{\theta}| > \varepsilon$ $c_3 > 0$; $\varepsilon > 0$

$= 0$; if $|\theta - \hat{\theta}| \leq \varepsilon$ (Typically, $c_3 = 1$)

- For instance, under Quadratic Loss, we can obtain the Bayes' estimator of β as:

$$\hat{\beta} = E[\beta|\boldsymbol{c}] = \int_0^1 \beta\, p(\beta|\boldsymbol{c}) d\beta$$

$$= \int_0^1 k\beta^{127.2}(1-\beta)^{30.8} \left\{\sum_{t=1}^{n}(c_t - \beta y_t)^2\right\}^{-n/2} d\beta$$

- The "posterior uncertainty" about β can be measured, for instance, by computing

$$var.(\beta|\boldsymbol{c}) = \int_0^1 (\beta - \hat{\beta})^2\, p(\beta|\boldsymbol{c}) d\beta$$

$$= \int_0^1 k(\beta - \hat{\beta})^2 \beta^{126.2}(1-\beta)^{30.8} \left\{\sum_{t=1}^{n}(c_t - \beta y_t)^2\right\}^{-n/2} d\beta$$

Numerical integration is again straightforward in each case.

- Similarly, if we wanted to draw inferences about σ, we would *marginalize* the Joint Posterior for both of the parameters by integrating out β.
- Then we would have to *normalize* the Marginal Posterior for σ, and we could look at the features of this posterior (*e.g.*, mean and variance) to draw inferences about this parameter.
- Specifically:

$$p(\sigma|\boldsymbol{c}) = \int_0^1 p(\beta, \sigma|\boldsymbol{c})d\beta$$

$$\propto \int_0^1 \sigma^{-(n+1)}\beta^{126.2}(1-\beta)^{30.8}\, exp\left\{-\frac{1}{2\sigma^2}\sum_{t=1}^{n}(c_t - \beta y_t)^2\right\}d\beta$$

- This integration would have to be done *numerically*.
- Then the *normalizing constant* for this Marginal Posterior is:

$$k' = 1/\int_0^{\infty}\int_0^{1} \sigma^{-(n+1)}\beta^{126.2}(1-\beta)^{30.8}\, exp\left\{-\frac{1}{2\sigma^2}\sum_{t=1}^{n}(c_t-\beta y_t)^2\right\}d\beta d\sigma$$

(*bivariate numerical integration* needed)

and the Marginal Posterior for σ would be

$$p(\sigma|\boldsymbol{c}) = k'\int_0^{1} \sigma^{-(n+1)}\beta^{126.2}(1-\beta)^{30.8}\, exp\left\{-\frac{1}{2\sigma^2}\sum_{t=1}^{n}(c_t-\beta y_t)^2\right\}d\beta$$

- This last integral is simple to evaluate numerically.
- Then, under Quadratic Loss (for example):

$$\hat{\sigma} = E[\sigma|\boldsymbol{c}] = \int_0^{\infty} \sigma\, p(\sigma|\boldsymbol{c})d\sigma$$

$$V[\sigma|\boldsymbol{c}] = \int_0^{\infty} (\sigma-\hat{\sigma})^2\, p(\sigma|\boldsymbol{c})d\sigma$$

R Code for Estimation of Consumption Function

```
cons.df<- read.table("http://web.uvic.ca/~dgiles/blog/consump.dat",header=T)

consump<- (cons.df$CONS-mean(cons.df$CONS))/1000

income<-  (cons.df$Y-mean(cons.df$Y)) /1000

year<- cons.df$YEAR

mle<- lm(consump~income -1)

c2<-consump^2

y2<-income^2

cy<-consump*income

pm<- 0.8

pv<- 0.001

n<- length(consump)

alpha<- (1/pv)*pm*(pm*(1-pm)-pv)
```

```
gamma<- (1/pv)*(1-pm)*(pm*(1-pm)-pv)

integrand<- function(b){(b^(alpha-1))*((1-b)^(gamma-1))/(sum(c2)+b^2*sum(y2)-
2*b*sum(cy))^(n/2)}

norm<- 1/integrate(integrand, lower = 0, upper = 1)$value

integrand<-  function(b) {b*(b^(alpha-1))*((1-b)^(gamma-1))
/(sum(c2)+b^2*sum(y2)-2*b*sum(cy))^(n/2)}

post_mean<- norm*integrate(integrand, lower = 0, upper = 1)$value

integrand<-  function(b) {(b-post_mean)^2*(b^(alpha-1))*((1-b)^(gamma-1))
/(sum(c2)+b^2*sum(y2)-2*b*sum(cy))^(n/2)}

post_var<- norm*integrate(integrand, lower = 0, upper = 1)$value

b<- seq(0.68 , 0.93, 0.00001)

prior<- dbeta(b,alpha, gamma)

sig<- summary(mle)$sigma

bmle<- summary(mle)$coef[1]
```

```
integrand<- function(b) {((n-1)*sig^2+(b-bmle)^2*sum(y2))^(-(n-1)/2)}    norml<-
1/integrate(integrand, lower = 0, upper = 1)$value

mlf<- norml*((n-1)*sig^2+(b-bmle)^2*sum(y2))^(-(n-1)/2)

posterior<- norm*((b)^(alpha-1))*((1-b)^(gamma-1))/(sum(c2)+b^2*sum(y2)

-2*b*sum(cy))^(n/2)

o<- order(posterior)

temp<- b[o]

post_mode<- temp[25001]

plot(b, posterior, type="l", lwd=2, col="blue",xlab="Beta", ylab="Densities")

lines(b, prior, lwd=2,col="red")

lines(b, mlf, lwd=2, col="black")

title(main="Consumption Function Example")

mtext("(Prior Mean = 0.8 ; Prior Variance = 0.001)")
```

```r
legend(0.7,50,lty=c(1,1,1), lwd=c(2,2,2), col=c("red","black","blue"),
c("Prior","Likelihood","Posterior"))

bmle                      # MLE of Beta
post_mean                  # Posterior Mean for Beta
post_var                  # Posterior Variance for Beta
post_mode                 # Posterior Mode for Beta
```

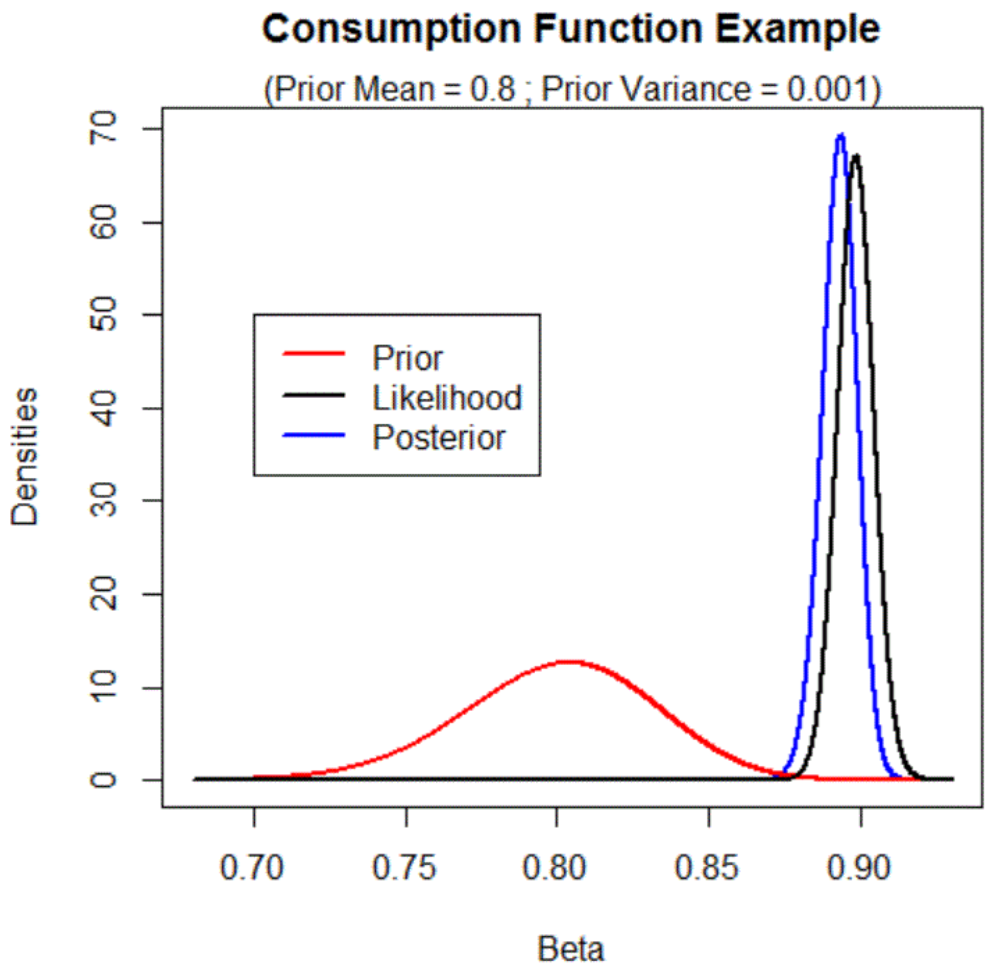

Posterior Mean = 0.893; Posterior Variance = $1.115*10^{-5}$; MLE = 0.898