STAT 703/J703

January 20th, 2005

-Lecture 4-

Instructor: Brian Habing

Department of Statistics

LeConte 203

Telephone: 803-777-3578

E-mail: habing@stat.sc.edu

Today

Sect 8.4: Method of Moments (cont).

- The Parametric Bootstrap
- Consistency

Sect 8.5: Maximum Likelihood

Given a MoM estimate, we still need to investigate its sampling distribution.

We can use the statistic to generate "new" bootstrap samples.

And then we can calculate the MoM estimator for each one of these. The set of these is the bootstrap distribution.

Ideally, the relationship between the *statistic* and *bootstrap distribution* should approximate the relationship between the *parameter* and the *sampling distribution*.


```
hist(lhat.dist,nclass=50,
  xlim=c(0,25))
lines(c(mean(lhat.dist),
  mean(lhat.dist)),
  c(-1000,30000),lwd=4,lty=5)
text(12,18000,"Bootstrap Mean")
lines(c(lhat,lhat),c(-
  1000,30000),lwd=4)
text(3,18000,"Original Est.")
```


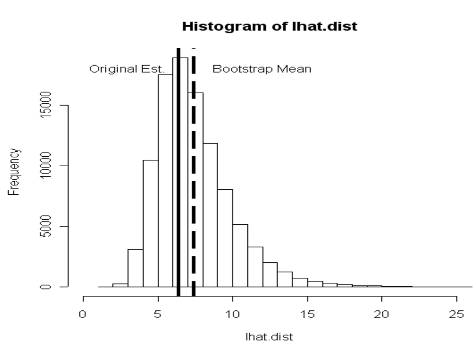

The MoM estimates do not necessarily correspond to a distribution that is likely to have produced them!!!!

Besides simplicity, there is another good property of MoM estimates however.

Definition: Let $\hat{\theta}_n$ be an estimate of θ for a sample of size n. $\hat{\theta}_n$ Is said to be a consistent estimator of θ if it converges to θ in probability.

That is, if for any $\varepsilon>0$,

$$P(|\hat{\theta}_n - \theta| > \varepsilon) \to 0 \text{ as } n \to \infty$$

8.5 – Maximum Likelihood

The idea behind maximum likelihood estimation is to find the parameters that seem most likely to have resulted in the observed statistic.

In the case of more than one observation, the likelihood is:

$\text{lik}(\theta|x_1,\ldots x_n)=f(x_1,\ldots x_n|\theta)=\Pi f(x_i|\theta)$

It is unpleasant to find the maximum of a product though…

An easier function to work with is the log likelihood

$L(q)=\log(\text{lik}(\theta))$

$=\log(\Pi f(x_i|\theta))$

$=\Sigma \log(f(x_i|\theta))$

Consider a sample from a Poisson distribution….

Unlike the Poisson distribution, a Gamma distribution has two parameters to deal with…


```
ngamloglike<-function(pars,data){
 a<-pars[1]
 l<-pars[2]
 n<-length(data)
 -1*
 (n*a*log(l)+(a1)*sum(log(data))-
  l*sum(data)-n*lgamma(a))
  }
optim(c(6,6),neggloglike,data=a)
```