Stat 342 Example 6

Consider a toy inference problem where

$$x \sim U(0,\theta)$$

so that

$$f(x|\theta) = \frac{1}{\theta} I[0<x<\theta]$$

Think of this as a density for x at a particular θ one has

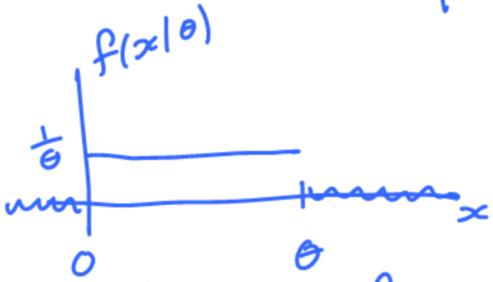

The likelihood function for an observed x looks like

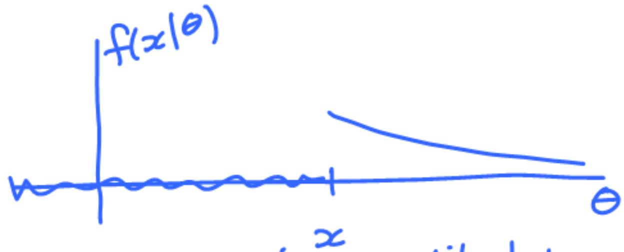

So the maximizer of the likelihood is

$$\hat{\theta}^{MLE}(x) = x$$

This estimator has SEL risk function

$$\text{MSE} \quad (R(\theta) = E_\theta (x-\theta)^2 = \text{Var}_\theta + \underbrace{(E_\theta x - \theta)^2}_{\text{squared bias}}$$

$$= \frac{\theta^2}{12} + \left(\frac{\theta}{2} - \theta\right)^2 = \frac{\theta^2}{3}$$

For sake of example, we can turn this into a Bayes inference problem by adopting a prior dsn for θ. Here we'll consider a $U(0,1)$ prior for θ, i.e. one with pdf

$$g(\theta) = I[0<\theta<1]$$

Then, a joint pdf for (x,θ) is

$$f(x|\theta)g(\theta) = \frac{1}{\theta} I[0<\theta<1]\, I[0<x<\theta]$$

$$= \begin{cases} \frac{1}{\theta} & 0 < x < \theta < 1 \\ 0 & \text{otherwise} \end{cases}$$

This is

Then given an observed x, the conditional density of θ (the "posterior" density $g(\theta|x)$) is proportional to

$$\frac{1}{\theta} I[x < \theta < 1]$$

(One fixes x and slices the picture above at that x.) Then

$$g(\theta|x) = \frac{\frac{1}{\theta} I[x<\theta<1]}{\int_x^1 \frac{1}{\theta}\, d\theta}$$

$$= \frac{1}{\theta} I[x<\theta<1] \Big/ (-\ln x)$$

So, e.g. the conditional (posterior) mean of θ given x is

$$\int \theta\, g(\theta|x)\, d\theta = \int_x^1 \theta \left(\frac{1}{\theta(-\ln x)}\right) dx = \frac{1-x}{(-\ln x)}$$

So, the SEL (Bayes) optimal estimator of θ based on x is

$$\hat{\theta}^{opt}(x) = \frac{1-x}{(-\ln x)}$$

The (Bayes) (single number) SEL risk of this estimator is

$$E\left(\left(\frac{1-x}{-\ln x}\right) - \theta\right)^2$$

this expected value averages out both x and θ using their joint distribution!

One thing we know about this expectation is for example that it is no larger than

$$E\left(R(\theta) \text{ for the MLE}\right) = E\,\frac{\theta^2}{3}$$

$$= \int_0^1 \frac{\theta^2}{3}\,d\theta = \frac{1}{9}$$

This averages out over θ according to its $U(0,1)$ marginal density

One may also take a Bayesian approach to inference about θ without employing decision analysis. For example, using the posterior density

$$g(\theta|x) = \frac{1}{\theta(-\ln x)}\, I\left[x<\theta<1\right]$$

one could set

$$.05 = \int_x^l \frac{1}{\theta(-\ln x)}\,d\theta \implies l = x^{.95}$$

and $$.05 = \int_u^1 \frac{1}{\theta(-\ln x)}\,d\theta \implies u = x^{.95}$$

and use the interval $(x^{.95}, x^{.05})$ to locate θ with 90% posterior probability. This is a 90% credible interval for θ (based on the $U(0,1)$ prior).

Note: $(x^{.95}, x^{.05})$ is not a 90% confidence interval for θ.

$$\text{confidence level}(\theta) = \text{coverage probability}(\theta)$$

$$= P_\theta\left[x^{.95} < \theta < x^{.05}\right]$$

$$= P_\theta\left[x > \theta^{20} \text{ and } x < \theta^{\frac{1}{.95}}\right]$$

$$= \frac{1}{\theta}\left(\theta^{.95} - \theta^{20}\right)$$

A plot of this shows that it's not $\geq .90$ even for $\theta \in [0,1]$.

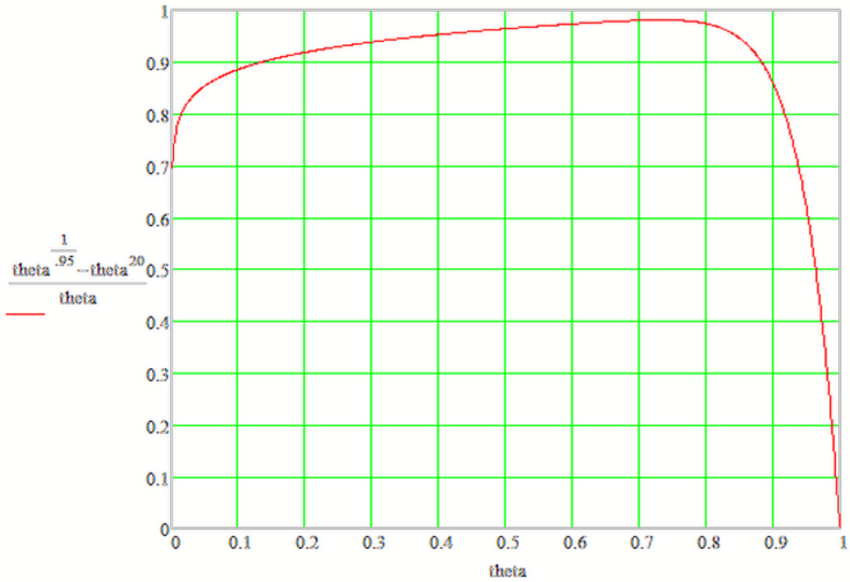

1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
theta
$\frac{\theta^{\frac{1}{.95}} - \theta^{20}}{\theta}$