and the size of the test is

$$P_{\mu=7}\left[\phi(X)=1\right] = 1-\left(\exp\left(-\frac{a(k'')}{7}\right)-\exp\left(-\frac{b(k'')}{7}\right)\right)$$

and by varying k'' we come up with a test of any desired size α

A bit on Set Estimation and Prediction

Set Estimation

$X \sim f(x|\theta) \quad \theta \in \Theta$

$\gamma(\theta)$ a parametric function of interest

Suppose for each x, $S(x)$ is a subset of $\gamma(\Theta)$ - We call $S(X)$ a (random) set estimator of $\gamma(\theta)$

Set Prediction $(X,Y) \sim f(x,y|\theta) \quad \theta \in \Theta$

Suppose that for each x, $S(x)$ is a subset of $\mathcal{Y}$ - We call $S(X)$ a set predictor of Y

Bayes Theory $\quad \theta \sim g(\theta)$

posterior dsn $g(\theta|x) \Rightarrow$ posterior dsn for $\gamma(\theta)$

where the posterior dsn of $\gamma(\theta)$ is cont's, it is common to look for a Highest Posterior Density set for $\gamma(\theta)$

day 32

e.g.

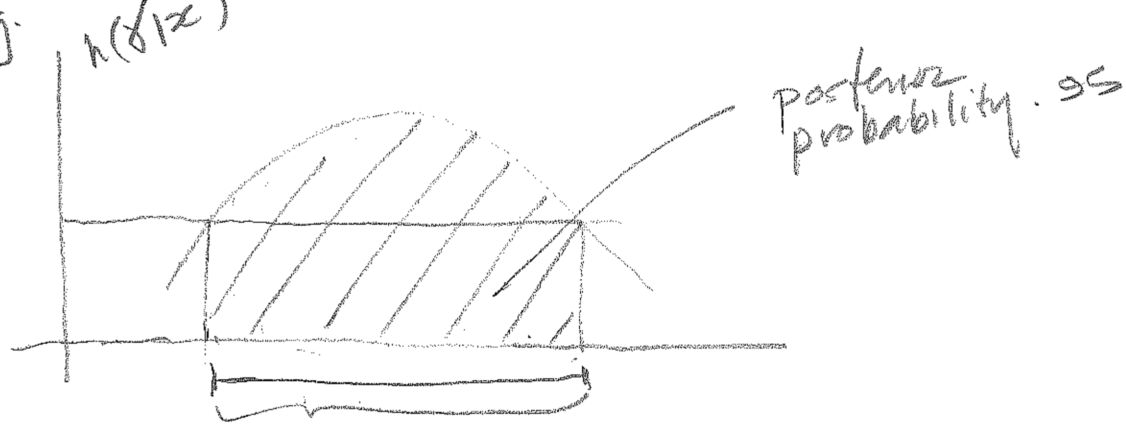

$S(x)$ — a 95% HPD "credible" interval for γ

Similarly, $\theta \sim g(\theta)$

$$f(x,y,\theta) = f(x,y|\theta)\, g(\theta)$$

$$f(x,y) = \int f(x,y,\theta)\, d\theta$$

This is the posterior predictive dsn of y given x $\rightarrow$ $f(y|x) = f(x,y)/f(x)$

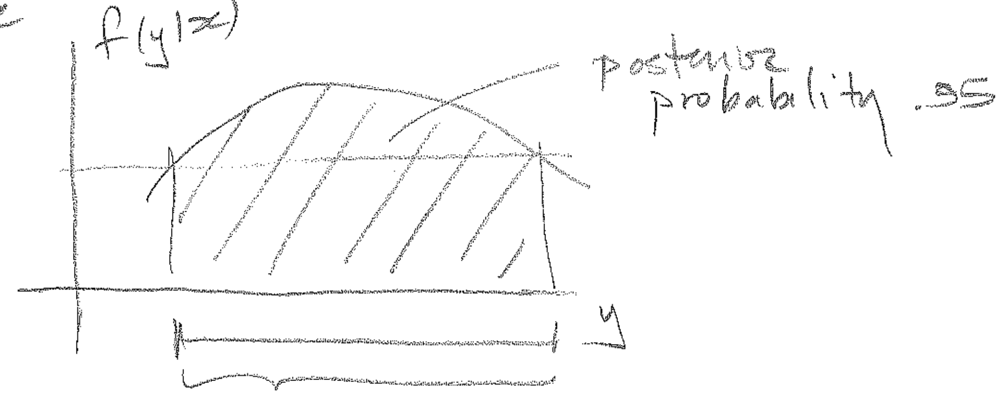

$S(x)$ — a 95% HPD "credible" interval for Y

<u>Example</u> $X_1, X_2, \ldots, X_n, X_{n+1}$ iid $N(\mu, 1)$

$\mu \sim N(0, \gamma^2)$ prior

Let $\bar{X} = \frac{1}{n}\sum_{i=1}^{n} X_i$

$$\mu \mid X_1, X_2, \ldots, X_n \text{ is } N\left[\left(\frac{n\gamma^2}{n\gamma^2+1}\right)\bar{X},\ \frac{\gamma^2}{n\gamma^2+1}\right]$$

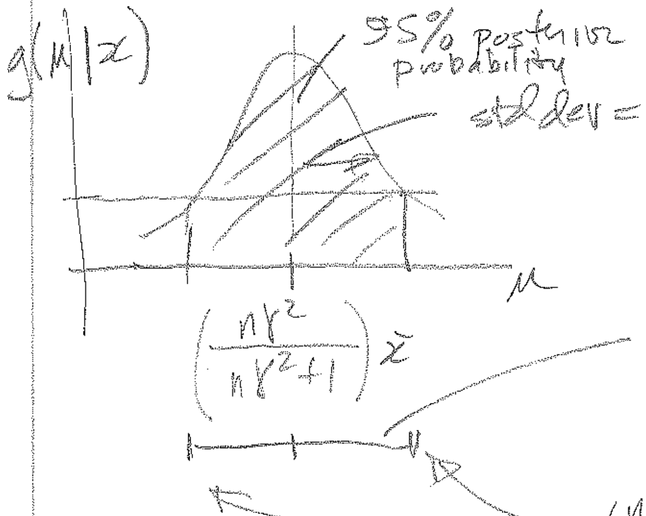

a 95% HPD credible interval for μ

$$\left(\frac{n\gamma^2}{n\gamma^2+1}\right)x \pm 1.96\sqrt{\left(\frac{\gamma^2}{n\gamma^2+1}\right)}$$

<u>Also</u>

$$X_{n+1} \mid X_1, \ldots, X_n \text{ is } N\left[\left(\frac{n\gamma^2}{n\gamma^2+1}\right)\bar{X},\ \underbrace{1+\frac{\gamma^2}{n\gamma^2+1}}_{\frac{(n+1)\gamma^2+1}{n\gamma^2+1}}\right]$$

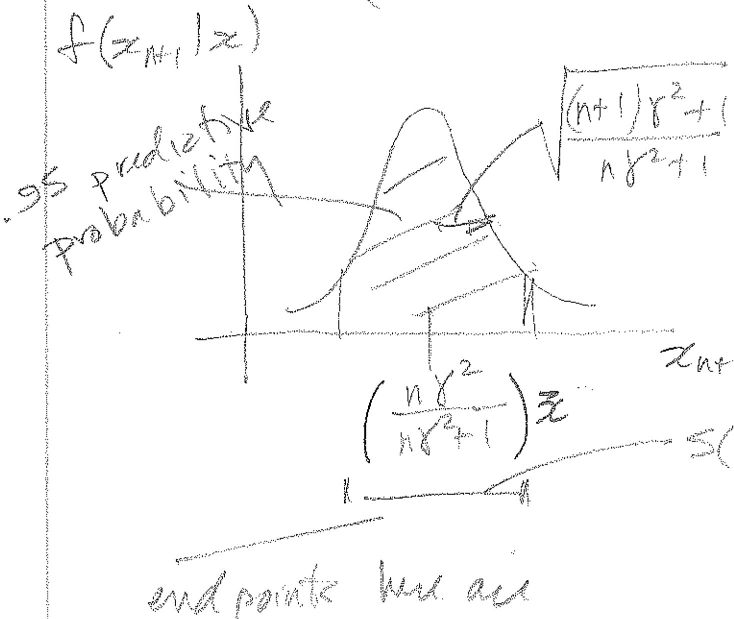

$S(x)$ a 95% HPD credible set for X_{n+1}

end points here are

$$\left(\frac{n\gamma^2}{n\gamma^2+1}\right)x \pm 1.96\sqrt{1+\frac{\gamma^2}{n\gamma^2+1}}$$

The theory of what to do to make Bayes set estimates and predictions is transparent – it's implementation may be not so easy – one relies on

MCMC to make it feasible to implement — how to identify the highest posterior density set is often not so simple (especially outside of the 1-D situation) — often one instead simply uses intervals (for 1-d γ's) that are of the form

$$\left(\begin{array}{c} \text{lower } \alpha/2 \\ \text{point of posterior} \\ \text{of } \gamma \end{array} , \begin{array}{c} \text{upper } \alpha/2 \\ \text{point of} \\ \text{posterior of } \gamma \end{array} \right)$$

as $1-\alpha$ level credible sets (without worrying about the "HPD" (and thus "smallest possible") form —

Frequentist Set Estimation and Prediction

Def The confidence level of $S(X)$ for $\gamma(\theta)$ is

$$\inf_\theta P_\theta\left[\gamma(\theta) \in S(X)\right]$$

Def The confidence level of $S(X)$ for Y is

$$\inf_\theta P_\theta\left[Y \in S(X)\right]$$

Standard Stat 500, 510 confidence set methods are often based on exact or approximate "pivots"/"pivotal quantities" — This is a function

$$p(X, \theta)$$

whose dsn is (at least approximately) free of θ i.e. the same for all θ

Example 1 sample normal

$$\frac{\bar{X}-\mu}{\frac{S}{\sqrt{n}}} \sim t_{n-1} \quad \text{produces the } t \text{ interval for } \mu$$

$\frac{(n-1)S^2}{\sigma^2} \sim \chi^2_{n-1}$ produces χ^2 interval for σ^2

2 - sample normal model

$\frac{S_1^2/\sigma_1^2}{S_2^2/\sigma_2^2} \sim F_{n_1-1,\, n_2-1}$ produces the F interval for $\frac{\sigma_1^2}{\sigma_2^2}$

Beyond normal linear models this kind of thing is pretty much case by case stuff – the only general theory for pivots is asymptotic / large sample theory (see Ch 5 upcoming)

Before going to Ch 5, there is one issue to discuss, namely the "duality" between tests and confidence sets

day 33

Still $X \sim f(x|\theta)$ with parameter space Θ

Suppose that for each $\theta_0 \in \Theta$

$\phi_{\theta_0}(x)$ is a size α test of $H_0: \theta = \theta_0$

Define

$$S(x) = \underbrace{\{\theta \in \Theta \mid \phi_\theta(x) < 1\}}$$

The set of parameters whose associated tests have positive (conditional on x) probability of acceptance

(if all of the ϕ_θ are non-randomized, $S(x)$ is the set of parameters whose associated tests accept)

Then $$P_\theta\left[\theta \in S(X)\right] = P_\theta\left[\phi_\theta(X) < 1\right]$$
$$= P_\theta\left[1 - \phi_\theta(X) > 0\right]$$
$$= E_\theta I\left[1 - \phi_\theta(X) > 0\right]$$
$$\geq E_\theta\left(1 - \phi_\theta(X)\right)$$
$$= 1 - \alpha$$

Complementary Fact: Confidence procedures can be used to do testing — Suppose that $S(x)$ is a set estimator of θ with

$$P_\theta\left[\theta \in S(X)\right] \geq \eta$$

Then define

$$\phi_\theta(x) = I\left[\theta \notin S(x)\right]$$

$$\pi_{\phi_\theta}(\theta) = P_\theta\left[\phi_\theta(X) = 1\right] = P_\theta\left[\theta \notin S(X)\right]$$
$$\leq 1 - \eta$$

and $\phi_{\theta_0}(x)$ is level (size) $\leq 1 - \eta$ for $H_0: \theta = \theta_0$

Example $X_1, X_2, \ldots X_n$ iid $N(\mu, 1)$

UMP size $\alpha = .05$ tests of $H_0: \mu = \mu_0$ vs $H_a: \mu > \mu_0$ are of the form

$$\phi_{\mu_0}(x) = \begin{cases} 1 & \text{if } \bar{x} > \mu_0 + 1.645\frac{1}{\sqrt{n}} \\ 0 & \text{otherwise} \end{cases}$$

and the corresponding 95% level confidence sets are

$$S(x) = \{\mu \mid \bar{x} \leq \mu + 1.645/\sqrt{n}\}$$
$$= \{\mu \mid \mu \geq \bar{x} - 1.645/\sqrt{n}\} = [\bar{x} - 1.645/\sqrt{n}, \infty)$$