

Recap:

- ① Cramer-Rao lower bound.
- ② Multivariate Cramer-Rao lower bound.

Methods of finding estimators

- (i) Method of moments
- (ii) Method of maximum likelihood estimation
- (iii) Bayes and minimax estimators

Method of moments

Sometimes we do not know how to construct estimators for a bunch of parameters and we need to depend on intuitions. In that case, creating a sample analogue of the population parameters might work as an estimator.

In general we create the method of moment estimator as following.

$$\text{Let } m_j = \frac{1}{n} \sum_{i=1}^n x_i^j,$$

$$\text{also assume } \mu'_j = E[x^j]$$

Generally μ'_j 's are functions of unknown parameters $\theta_1, \dots, \theta_R$. We find ~~at~~ the method of moment estimators (MOM) of $\theta_1, \dots, \theta_R$.

by solving K equations

$$m_j = \mu_j(\theta_1, \dots, \theta_K), j=1, \dots, K.$$

Example: $x_1, \dots, x_n \stackrel{iid}{\sim} N(\mu, \sigma^2)$

We will use, $m_1 = E[x]$

$$m_2 = E[x^2]$$

$$\Rightarrow \frac{1}{n} \sum_{i=1}^n x_i = E[x] \quad \dots \quad (1)$$

$$\bullet \quad \frac{1}{n} \sum_{i=1}^n x_i^2 = E[x^2] \quad \dots \quad (2)$$

$$\Rightarrow \frac{1}{n} \sum_{i=1}^n x_i = \mu$$

$$\frac{1}{n} \sum_{i=1}^n x_i^2 = \mu^2 + \sigma^2 \Rightarrow \hat{\mu} = \frac{1}{n} \sum_{i=1}^n x_i, \\ \hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n x_i^2 - \bar{x}^2 \\ = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2$$

In this case, MOM estimators are functions of complete, sufficient statistics. However, it is not the case in many distributions.

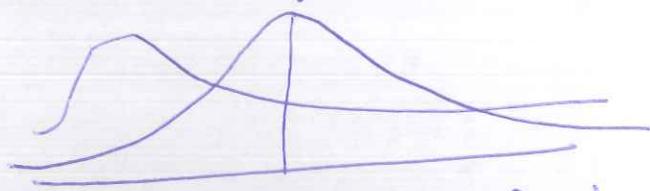
e.g. $x_1, \dots, x_n \stackrel{iid}{\sim} U(\mu-2, \mu+2)$

$$\frac{1}{n} \sum_{i=1}^n x_i = E[x] = \mu$$

\Rightarrow MOM estimator of μ is \bar{x} .

Here minimal sufficient statistic is a function of $x_{(1)}, x_{(n)}$.. Hence MOM estimator is not

a function of ~~the~~ minimal sufficient statistic.



Another problem of the MOM estimator is that it ~~might~~ not ~~provide~~ provide estimators in the support of the parameter space.

Example: $x_1, \dots, x_n \stackrel{iid}{\sim} \text{Bin}(K, p)$, K and p are both unknown.

$$\frac{1}{n} \sum_{i=1}^n x_i = E[X] \Rightarrow \frac{1}{n} \sum_{i=1}^n x_i = Kp$$

$$\frac{1}{n} \sum_{i=1}^n x_i^2 = E[X^2] \quad \frac{1}{n} \sum_{i=1}^n x_i^2 = Kp(1-p) + K^2 p^2$$

$$\Rightarrow p = \frac{\bar{x}}{K} \quad \text{and} \quad \frac{1}{n} \sum_{i=1}^n x_i^2 = Kp - Kp^2 + K^2 p^2 \\ = \bar{x} - \cancel{Kp} + Kp^2(K-1)$$

$$\Rightarrow \frac{1}{n} \sum_{i=1}^n x_i^2 - \bar{x} = K \left(\frac{\bar{x}}{K} \right)^2 (K-1)$$

$$\Rightarrow \frac{1}{n} \sum_{i=1}^n x_i^2 - \bar{x} = \frac{\bar{x}^2}{K} (K-1)$$

$$\Rightarrow K \left[\frac{1}{n} \sum_{i=1}^n x_i^2 - \bar{x} \right] = \bar{x}^2 K - \bar{x}^2$$

$$\Rightarrow K = \frac{\bar{x}^2}{\bar{x}^2 - \frac{1}{n} \sum_{i=1}^n x_i^2 + \bar{x}}$$

the R.H.S. may not be an integer while L.H.S. is an integer. Hence MOM estimator is not even a feasible estimator here.

Example: $x_1, \dots, x_n \stackrel{iid}{\sim} N(\theta, \theta^2)$, $\theta > 0$

$$\frac{1}{n} \sum_{i=1}^n x_i = \bar{x}$$

and $\frac{1}{n} \sum_{i=1}^n x_i^2 = E[X^2] = \theta^2 + \theta^2 = 2\theta^2$
 $\Rightarrow \bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$ and $\theta = \sqrt{\frac{1}{2n} \sum_{i=1}^n x_i^2}$

Maximum Likelihood estimator

Suppose $f_{\theta_1, \dots, \theta_K}(x)$ be the p.d.f. of X . Then the likelihood of $\underline{x} = (x_1, \dots, x_n)$ is given by

$$L(\theta_1, \dots, \theta_K) = \prod_{i=1}^n f_{\theta_1, \dots, \theta_K}(x_i)$$

Maximum Likelihood estimator (MLE) is the value of $\hat{\theta} = (\hat{\theta}_1, \dots, \hat{\theta}_K)$ that maximizes the likelihood $L(\theta_1, \dots, \theta_K)$. MLE is found by

solving $\frac{\partial L(\theta)}{\partial \theta_i} = 0, i=1, \dots, K$.

sometimes you use $\frac{\partial \log L(\theta)}{\partial \theta_i} = 0, i=1, \dots, K$.

to find the MLE.

Note that $\frac{\partial \log L(\theta)}{\partial \theta_i} = 0 \Rightarrow \frac{1}{L(\theta)} \frac{\partial L(\theta)}{\partial \theta_i} = 0$

$\Rightarrow \frac{\partial L(\theta)}{\partial \theta_i} = 0$. Thus they are equivalent.

$$\frac{\partial \log L(\theta)}{\partial \theta_i} = 0 \Rightarrow \frac{\partial \left[\sum_{j=1}^n \log f_{\theta_1, \dots, \theta_K}(x_j) \right]}{\partial \theta_i} = 0$$

$$\Rightarrow \sum_{j=1}^n \frac{\partial \log f_{\theta_1, \dots, \theta_K}(x_j)}{\partial \theta_i} = 0, i=1, \dots, K.$$

Note that $\frac{\partial \log f_{\theta_1, \dots, \theta_k}(x_j)}{\partial \theta_i}$ for $i=1 \dots, k$
are components of the score vector.

In the one dimensional case where $k=1$,
the equations boil down to

$$\sum_{j=1}^n \frac{\partial \log f_\theta(x_j)}{\partial \theta} = 0 \Rightarrow \sum_{j=1}^n u_\theta(x_j) = 0$$

Let $\hat{\theta}_1, \dots, \hat{\theta}_k$ be the solutions of the system
of equations to find the MLE. The MLE should
also satisfy that $\frac{\partial^2 L(\theta)}{\partial \theta^2} \Big|_{\theta=\hat{\theta}} < 0$

MLE sometimes lies in the boundary of the
parameter space.

Example: $x_1, \dots, x_n \stackrel{iid}{\sim} N(\theta, 1)$, $\theta > 0$.

Find the MLE of θ .

When $\theta > 0$ restriction is not there, MLE
of θ is given by \bar{x} . However, when we
have this additional restriction we
can't use \bar{x} as the MLE.

If $\bar{x} > 0$, then obviously \bar{x} is going
to be the MLE.

If $\bar{x} \leq 0$ then \bar{x} is not the MLE.

$$\begin{aligned} L(\theta) &= \prod_{i=1}^n \frac{1}{\sqrt{2\pi}} \exp \left\{ -\frac{(x_i - \theta)^2}{2} \right\} \\ &= \left(\frac{1}{\sqrt{2\pi}} \right)^n \exp \left\{ -\sum_{i=1}^n \frac{(x_i - \theta)^2}{2} \right\} \\ &= \left(\frac{1}{\sqrt{2\pi}} \right)^n \exp \left\{ -\frac{1}{2} \left[\sum_{i=1}^n (x_i - \bar{x} + \bar{x} - \theta)^2 \right] \right\} \\ &= \left(\frac{1}{\sqrt{2\pi}} \right)^n \exp \left\{ -\frac{1}{2} \left[\sum_{i=1}^n (x_i - \bar{x})^2 + (\bar{x} - \theta)^2 + 2(x_i - \bar{x})(\bar{x} - \theta) \right] \right\} \\ &= \underbrace{\left(\frac{1}{\sqrt{2\pi}} \right)^n \exp \left\{ -\frac{1}{2} \sum_{i=1}^n (x_i - \bar{x})^2 \right\}}_{\text{exp}\left\{ -\frac{n}{2} (\bar{x} - \theta)^2 \right\}} \end{aligned}$$

Note that $\bar{x} \leq 0$ so, let $\bar{x} = -z$, $z \geq 0$

$$\begin{aligned} L(\theta) &\propto \exp \left\{ -\frac{n}{2} (-z - \theta)^2 \right\} \\ &= \exp \left\{ -\frac{n}{2} (z + \theta)^2 \right\} \end{aligned}$$

Note that the range of θ is $\theta > 0$.

As θ increases the function $L(\theta)$ monotonically decreases.

The maximum is achieved when $\theta = 0$ which is outside the range of θ .
Thus no MLE exists in the case where $\bar{x} \leq 0$.

However if we begin with the distribution $x_1, \dots, x_n \stackrel{\text{iid}}{\sim} N(\theta, 1)$, $\theta \geq 0$
 then when $\bar{x} \geq 0$ $\theta = \bar{x}$ maximizes the likelihood.

When $\bar{x} < 0$, the likelihood is maximized at $\theta = 0$.

MLE of θ in this problem is $\max\{0, \bar{x}\}$

Remark: MLE does not need to be a function of ~~a~~ minimal sufficient statistic.

$x_1, \dots, x_n \stackrel{\text{iid}}{\sim} U(\theta-1, \theta+1)$

$$L(\theta) = \frac{1}{2^n} I[x_{(n)} - 1 < \theta < x_{(1)} + 1]$$

Any value of $\theta \in (x_{(n)} - 1, x_{(1)} + 1)$ maximizes the likelihood. Hence MLE is not unique.

Also,

$$\hat{\theta} = \left(\frac{\bar{x}}{\bar{x}+1}\right)(x_{(n)} - 1) + \frac{1}{(\bar{x}+1)}(x_{(1)} + 1)$$

$\hat{\theta} \in (x_{(n)} - 1, x_{(1)} + 1)$ being ~~a~~ a weighted average of $x_{(n)} - 1$ and $x_{(1)} + 1$.

Thus $\hat{\theta}$ is an MLE, but it is not a function of a minimal sufficient statistic.

Invariance property of MLE

If $\hat{\theta}$ is an MLE of θ , $\tau(\hat{\theta})$ is an MLE of $\tau(\theta)$, for any function $\tau(\theta)$.

Recap: ① Method of moments estimator

② Maximum Likelihood estimator.

Invariance property of MLE:

If $\hat{\theta}$ is the MLE of θ then $\tau(\hat{\theta})$ is the MLE of $\tau(\theta)$.

Asymptotic distribution of MLE:

In smooth regular problems MLE $\hat{\theta}$ has the property that $\sqrt{n}(\hat{\theta} - \theta) \xrightarrow{d} N(0, \frac{1}{I(\theta)})$

③ If g is any function of θ s.t. $g'(\theta) \neq 0$.

$\sqrt{n}(g(\hat{\theta}) - g(\theta)) \xrightarrow{d} N\left(0, \frac{[g'(\theta)]^2}{I(\theta)}\right)$ (By Delta theorem).

when n is very large

the distribution of $\sqrt{n}(g(\hat{\theta}) - g(\theta))$ can be approximated by $N\left(0, \frac{[g'(\theta)]^2}{I(\theta)}\right)$

$$E[\sqrt{n}(g(\hat{\theta}) - g(\theta))] \approx 0 \Rightarrow E[g(\hat{\theta})] \approx g(\theta)$$

$$\text{and } \text{Var}(\sqrt{n}(g(\hat{\theta}) - g(\theta))) \approx \frac{[g'(\theta)]^2}{I(\theta)}$$

$$\Rightarrow \text{Var}(g(\hat{\theta})) \approx \frac{[g'(\theta)]^2}{n I(\theta)}$$

thus $g(\hat{\theta})$ is an asymptotically unbiased estimator of $g(\theta)$ whose variance asymptotically achieves the Cramér-Rao lower bound.

Asymptotic distribution of MLE is Normal
when the regularity conditions hold.

It is not always true.

Example: $x_1, \dots, x_n \stackrel{iid}{\sim} u(0, \theta)$, $\theta > 0$.

$$f(x_1, \dots, x_n) = \frac{1}{\theta^n} I(x_{(n)} < \theta)$$

$$\text{If } \theta > x_{(n)} \text{ then } \frac{1}{\theta} < \frac{1}{x_{(n)}}$$

Here the MLE is $x_{(n)}$.

$$\text{Now, } \sqrt{n}(x_{(n)} - \theta) \not\overset{d}{\rightarrow} N(0, \frac{1}{I(\theta)})$$

$$n(x_{(n)} - \theta) \xrightarrow{d} \text{Exp}(\theta)$$

Bayes and Minimax estimator

$$x_1, \dots, x_n \stackrel{iid}{\sim} f(x|\theta)$$

Assume a prior distribution on θ .

Suppose the prior dist. is $\pi(\theta)$. There are objective and subjective choices of priors.

Posterior dist., denoted by $\pi(\theta|x_1, \dots, x_n)$,

$$\pi(\theta|x_1, \dots, x_n) = \frac{f(x_1, \dots, x_n|\theta) \pi(\theta)}{\int f(x_1, \dots, x_n|\theta) \pi(\theta) d\theta}$$

Example: $x_1, \dots, x_n \stackrel{iid}{\sim} \text{Ber}(p) \Rightarrow p \sim \text{Beta}(\alpha, \beta)$

$$\pi(p|x_1, \dots, x_n) \propto \left[\prod_{i=1}^n p^{x_i} (1-p)^{1-x_i} \right] p^{\alpha-1} (1-p)^{\beta-1}$$

$$\propto p^{\sum_{i=1}^n x_i + \alpha - 1} (1-p)^{n - \sum_{i=1}^n x_i + \beta - 1}$$

Thus $p|x_1, \dots, x_n \sim \text{Beta}\left(\sum_{i=1}^n x_i + \alpha, n - \sum_{i=1}^n x_i + \beta\right)$

Example: $x_1, \dots, x_n \stackrel{iid}{\sim} \text{Poi}(\lambda), \lambda \sim \text{Gamma}(\alpha, \beta)$

$$\pi(\lambda|x_1, \dots, x_n) \propto \left[\prod_{i=1}^n \frac{e^{-\lambda} \lambda^{x_i}}{x_i!} \right] e^{-\lambda/\beta} \lambda^{\alpha-1}$$

$$\propto e^{-\lambda(n+\frac{1}{\beta})} \lambda^{\sum_{i=1}^n x_i + \alpha - 1}$$

$$\lambda|x_1, \dots, x_n \sim \text{Gamma}\left(\sum_{i=1}^n x_i + \alpha, \frac{1}{n + \frac{1}{\beta}}\right)$$

Definition (Conjugate family)

Let \mathcal{F} denote the class of p.d.f.s or p.m.f.s
 Let Π denote the class of prior distributions
 $f(x|\theta)$. A class $\Theta \subset \Pi$ of prior distributions
 is a conjugate family for \mathcal{F} if the
 posterior distribution is in the class Π
 for all $f \in \mathcal{F}$ and all priors in Π .

Now we will learn how to construct good
 estimators from Bayesian statistic.
 We will be using the "risk function"
 notion in Bayesian statistics.

Let $\delta(\underline{x})$ be an estimator of θ . The loss in estimating θ by $\delta(\underline{x})$ is represented by a function known as the loss function.

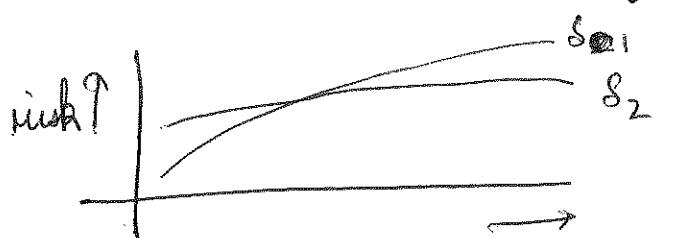
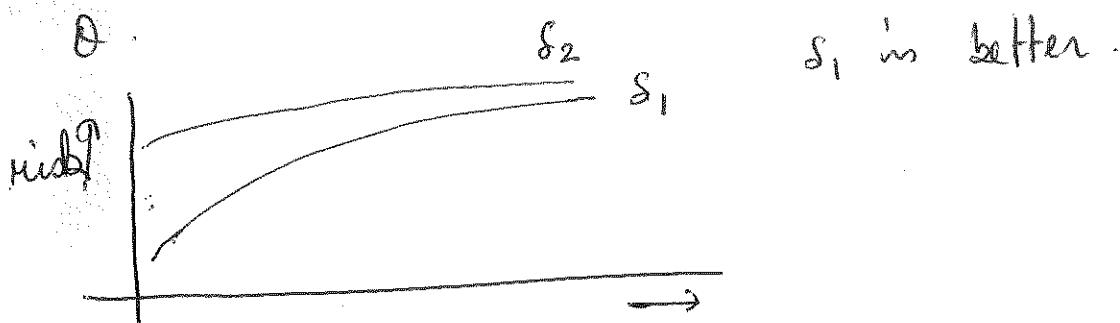
Mathematically we write it as $L(\theta, \delta)$

Define $R(\theta, \delta) = E_{\underline{X} \sim P} [L(\theta, \delta(\underline{x}))]$ is known as the risk function of δ .

when $L(\theta, \delta(\underline{x})) = (\delta(\underline{x}) - \theta)^2$ then

$R(\theta, \delta(\underline{x})) = E_{\underline{X} \sim P} [(\delta(\underline{x}) - \theta)^2]$ is the MSE of $\delta(\underline{x})$.

Given two estimators δ_1, δ_2 , we say δ_1 is better than δ_2 if $R(\theta, \delta_1) \leq R(\theta, \delta_2)$ for all θ and the inequality is strict for at least one



When two risks intersect each other, we do not know which estimator to choose. At that time, we have to propose a summary measure from the entire risk fn.

Different summary measures have been proposed.
We will discuss the following two.

$$1. \text{ Average risk} = E_{\theta} E_{X|D} [L(\theta, \delta(x))].$$

Minimizing this average risk will lead to the Bayes estimator.

$$2. \text{ Supremum risk} = \sup_{\theta} E_{X|D} [L(\theta, \delta(x))].$$

Minimizing the supremum risk will give rise to the Minimax estimator.

Bayes estimation

$$E_{\theta} E_{X|D} [L(\theta, \delta(x))] = E_{X, \theta} [L(\theta, \delta(x))] = E_X E_{\theta|X} [L(\theta, \delta(x))]$$

we have to minimize this Average risk w.r.t. $\delta(x)$.

we will find $\delta(x)$ that minimizes $E_{\theta|X} [L(\theta, \delta(x))]$ for all x .

There are a few loss functions where it can be done easily.

Squared error loss

$$\text{Let } L(\theta, \delta(x)) = (\delta(x) - \theta)^2$$

$$\begin{aligned} E_{\theta|X} [(\delta(x) - \theta)^2] &= E_{\theta|X} [(\delta(x) - E[\theta|x] + E[\theta|x] - \theta)^2] \\ &= E_{\theta|X} [(\delta(x) - E[\theta|x])^2] + E_{\theta|X} [(\theta - E[\theta|x])^2] \end{aligned}$$

$$E_{\theta|x}[(\delta(x) - \theta)^2] \geq E_{\theta|x}[(\cancel{E}[\theta|x] - \theta)^2]$$

thus $\delta(\underline{x}) = E[\theta|x]$ is going to be the Bayes estimator.

$E_{\theta|x}[L(\theta, \delta(\underline{x}))]$ you want to minimize w.r.t. $\delta(\underline{x})$

$$\frac{d}{ds} E_{\theta|x}[L(\theta, s)] = 0$$

~~$$\frac{d}{ds}$$~~ For $L(\theta, s) = (\theta - s)^2$

$$\frac{d}{ds} E_{\theta|x}[(\theta - s)^2] = 0$$

$$\Rightarrow E_{\theta|x}[2(s - \theta)] = 0$$

$$\Rightarrow E_{\theta|x}[s] = E_{\theta|x}[\theta]$$

$$\Rightarrow \delta(\underline{x}) = E_{\theta|x}[\theta]$$

Weighted Squared error loss:

$$L(\theta, \delta(\underline{x})) = \omega(\theta)(\theta - \delta(\underline{x}))^2, \omega(\theta) > 0$$

$$\frac{d}{ds} E_{\theta|x}[\omega(\theta)(\theta - s)^2] = 0$$

$$\Rightarrow E_{\theta|x}[2\omega(\theta)(s - \theta)] = 0$$

$$\Rightarrow E_{\theta|x}[\omega(\theta)s] = E_{\theta|x}[\theta\omega(\theta)]$$

$$\Rightarrow \delta(\underline{x}) = \frac{E_{\theta|x}[\theta\omega(\theta)]}{E_{\theta|x}[\cancel{\omega(\theta)}]} \quad \textcircled{6}$$

Example: $X_1, \dots, X_n \stackrel{iid}{\sim} \text{Ber}(p)$, $p \sim \text{Beta}(\alpha, \beta)$

$$p | X_1, \dots, X_n \sim \text{Beta}\left(\alpha + \sum_{i=1}^n X_i, \beta + n - \sum_{i=1}^n X_i\right)$$

under squared error loss $L(\theta, \delta(\underline{x})) = (\theta - \delta(\underline{x}))^2$
What is the Bayes estimator?

$$\text{Bayes estimator } E[p | X_1, \dots, X_n] = \frac{\alpha + \sum_{i=1}^n X_i}{\alpha + \sum_{i=1}^n X_i + n - \sum_{i=1}^n X_i + \beta}$$

$$= \frac{\alpha + \sum_{i=1}^n X_i}{\alpha + \beta + n}$$

$$E\left[E[p | X_1, \dots, X_n]\right] = \frac{\alpha + np}{\alpha + \beta + n} \neq p$$

this is a biased estimator of p .

Theorem: No unbiased estimator $\delta(\underline{x})$ of θ can be a Bayes estimator under squared error loss unless $E_\theta E_{\underline{x}|\theta}[(\theta - \delta(\underline{x}))^2] = 0$

Pf: Let $\delta(\underline{x})$ be an unbiased estimator which is also a Bayes estimator.

$$E[\delta(\underline{x}) | \theta] = \theta \text{ for all } \theta \text{ and } E[\theta | \underline{x}] = \delta(\underline{x})$$

Now,

$$E_\theta E_{\underline{x}|\theta} [\delta(\underline{x}) \theta] = E_\theta [\theta E_{\underline{x}|\theta} [\delta(\underline{x})]] = E_\theta [\theta]$$

~~$E_{\underline{x}} E_\theta [\theta]$~~

$$E_{\underline{x}} E_{\theta | \underline{x}} [\delta(\underline{x}) \theta] = E_{\underline{x}} [\delta(\underline{x}) E_{\theta | \underline{x}} [\theta]] = E_{\underline{x}} [\delta(\underline{x})]$$

$$E_{\theta} E_{x|D} [\delta(x)^\gamma] = E_{\theta} E_{x|D} [\delta(x)^\theta] = E_{\theta} E_{x|D} [\theta^\gamma] \quad \text{--- (*)}$$

Hence,

$$E_{\theta, x} [(\theta - \delta(x))^\gamma] = E_{\theta, x} [\theta^\gamma - 2\delta(x)\theta + \delta(x)^\gamma]$$

By (*) $= 0$

Sometimes you can show that the Bayes estimator is a convex combination of the prior mean and data mean.

Example: $x_1, \dots, x_n \stackrel{iid}{\sim} \text{Ber}(\theta)$, $\theta \sim \text{Beta}(\alpha, \beta)$

$$\begin{aligned} \Rightarrow E[\theta | x_1, \dots, x_n] &= \frac{\sum_{i=1}^n x_i + \alpha}{\alpha + \beta + n} \\ &= \underbrace{\frac{\sum_{i=1}^n x_i}{n}}_{\text{data mean}} \cdot \underbrace{\frac{n}{\alpha + \beta + n}}_{w_1} + \underbrace{\left(\frac{\alpha}{\alpha + \beta}\right)}_{\text{prior mean}} \underbrace{\frac{(\alpha + \beta)}{\alpha + \beta + n}}_{w_2} \end{aligned}$$

$$w_1 + w_2 = 1$$

Midterm

1. $x_1, \dots, x_{2n} \stackrel{iid}{\sim} \text{Ber}(\theta)$, $s_n = \sum_{i=1}^n x_{2i} x_{2i-1}$

$$\sqrt{n} \left(\frac{s_n}{n} - p^2 \right) \quad E\left[\frac{s_n}{n}\right] = p^2$$

$$\begin{aligned} \text{Var}(\cancel{x_{2i} x_{2i-1}}) &= E[x_{2i}^2 x_{2i-1}^2] - \{E[x_{2i}] E[x_{2i-1}]\}^2 \\ &= E[x_{2i}^2] E[x_{2i-1}^2] - p^4 \\ &= \{ \cancel{\cancel{p^2(1-p)^2}} + p^2 \} \{ p(1-p) + p^2 \} - p^4 \end{aligned}$$

$$\text{var}(X_2 | X_{2i-1}) = p^2 - p^4 = p^2(1-p^2)$$

CLT tells us $\sqrt{n} \left(\frac{S_n}{n} - p^2 \right) \xrightarrow{d} N(0, p^2(1-p^2))$

$$2. \quad X = \begin{cases} -1 & \text{w.p. } 0/2 \\ 0 & \text{w.p. } 1-0 \\ +1 & \text{w.p. } 0/2 \end{cases}$$

$$|X_i| = 1 \quad \text{w.p. } 0 \\ = 0 \quad \text{w.p. } 1-0$$

$\Rightarrow (\cancel{X}_i, |X_1|, \dots, |X_n|) \stackrel{\text{iid}}{\sim} \text{Ber}(0)$.

Now the rest follows from classes.

Or that was $\frac{\sum_{i=1}^n |X_i| \left(\sum_{i=1}^n |X_i| - 1 \right)}{n(n-1)}$

$$3. \quad X_1, \dots, X_n \stackrel{\text{iid}}{\sim} N(\mu, \sigma^2),$$

(\bar{X}, S^2) is a complete sufficient stat.

$E[\bar{X}] = \mu$ thus By Lehman-Scheffe theorem
 \bar{X} is the UMVUE for μ .

$$\frac{(n-1)S^2}{\sigma^2} \sim \chi_{n-1}^2 \quad \text{=} \text{Gamma}\left(\frac{n-1}{2}, 2\right)$$

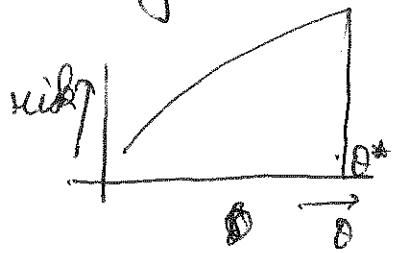
$$\Rightarrow \frac{\sigma^2}{(n-1)S^2} \sim \text{Inverse-Gamma}\left(\frac{n-1}{2}, 2\right)$$

$$\Rightarrow E\left[\frac{\sigma^2}{(n-1)S^2}\right] = \frac{1}{2\left(\frac{n-1}{2}-1\right)} = \frac{1}{n-3}$$

$$\Rightarrow E\left[\frac{n-3}{(n-1)S^2}\right] = \frac{1}{\sigma^2} \quad \text{this is a function of the complete sufficient stat.}$$

Recap:

- ① Risk function $E_{\underline{X}|\theta} [L(\theta, \delta(\underline{x}))]$ of an estimator $\delta(\underline{x})$.
- ② $E_\theta E_{\underline{X}|\theta} [L(\theta, \delta(\underline{x}))] \rightarrow$ an estimator $\delta(\underline{x})$ that minimizes this average risk fn. is called the Bayes estimator
- ③ $\sup_\theta E_{\underline{X}|\theta} [L(\theta, \delta(\underline{x}))]$. The Minimax estimator minimizes $\sup_\theta E_{\underline{X}|\theta} [L(\theta, \delta(\underline{x}))]$. Thus minimax estimator is going to protect us from the worst case scenario.
- ④ Computing/finding minimax estimator is difficult. However, we can use other tricks to compute minimax estimators. If the prior distribution of θ gives high probability to those values, which will increase the risk, then intuitively a minimax estimator is a Bayes estimator.



If the prior only has mass at θ^* then the

$$E_\theta E_{\underline{X}|\theta} [L(\theta, \delta(\underline{x}))] = \sup_\theta E_{\underline{X}|\theta} [L(\theta, \delta(\underline{x}))]$$

Definition (Least favorable distribution)

A prior distribution $\pi(\theta)$ is known to be a least favorable prior if $E_\theta E_{\underline{X}|\theta} [L(\theta, \delta_\pi(\underline{X}))] \geq E_\theta E_{\underline{X}|\theta} [L(\theta, \delta_{\pi'}(\underline{X}))]$ for all ~~prior~~ priors distributions $\pi'(\theta)$ on θ . Here s_π and $s_{\pi'}$ are Bayes estimators w.r.t. priors π and π' respectively.

So, $\pi(\theta)$ is such that ~~it~~ it increases the Bayes risk for the Bayes estimator w.r.t. any other prior.

Result: If $\pi(\theta)$ be a prior distribution for which $\int \underline{E}_\theta [L(\theta, s_\pi)] d\theta = \sup_\theta E_{\underline{X}|\theta} [L(\theta, s_\pi)]$

where s_π is the Bayes estimator. Then

(a) s_π is minimax estimator

(b) π is least favorable distribution

Proof: a) For any $s(\underline{X})$,

$$\sup_\theta E_{\underline{X}|\theta} [L(\theta, s(\underline{X}))] \geq E_\theta E_{\underline{X}|\theta} [L(\theta, s(\underline{X}))]$$

$$\geq E_\theta E_{\underline{X}|\theta} [L(\theta, s_\pi(\underline{X}))] = \sup_\theta E_{\underline{X}|\theta} [L(\theta, s_\pi)]$$

by assumption

thus $s_\pi(\underline{X})$ is a minimax estimator

(b) Note that, by assumption

$$E_0 E_{\underline{\pi}(\theta)} [L(\theta, \delta_{\pi})] = \sup_{\theta} E_{\underline{\pi}(\theta)} [L(\theta, \delta_{\pi})]$$

$$\geq \int E_{\underline{\pi}(\theta)} [L(\theta, \delta_{\pi})] \pi'(\theta) d\theta \quad \text{for any other priors dist. } \pi'(\theta) \text{ on } \theta$$

$$E_{\underline{\pi}(\theta)} [L(\theta, \delta_{\pi})] \leq \sup_{\theta} E_{\underline{\pi}(\theta)} [L(\theta, \delta_{\pi})]$$

$$\int E_{\underline{\pi}(\theta)} [L(\theta, \delta_{\pi})] \pi'(\theta) d\theta \leq \sup_{\theta} E_{\underline{\pi}(\theta)} [L(\theta, \delta_{\pi})] \int \pi'(\theta) d\theta$$

$$\geq \int E_{\underline{\pi}(\theta)} [L(\theta, \delta_{\pi'})] \pi'(\theta) d\theta$$

$$= E_0 E_{\underline{\pi}(\theta)} [L(\theta, \delta_{\pi'})]$$

thus, $E_0 E_{\underline{\pi}(\theta)} [L(\theta, \delta_{\pi})] \geq E_0 E_{\underline{\pi}(\theta)} [L(\theta, \delta_{\pi'})]$
 which implies that $\underline{\pi}(\theta)$ is the least favorable distribution.

To use this result, we need to have a prior dist. for which

$$\int E_{\underline{\pi}(\theta)} [L(\theta, \delta_{\pi})] \pi(\theta) d\theta = \sup_{\theta} E_{\underline{\pi}(\theta)} [L(\theta, \delta_{\pi})]$$

then we can find the Bayes estimator w.r.t. that prior. Our result says that this Bayes estimator will be the minimax estimator

If the risk fn. $E_{\bar{X}|\theta} [L(\theta, \delta_{\pi})]$ is const.
then average is going to be the supremum.

One algorithm

Step 1: Take a prior dist of your choice.

Step 2: Compute the risk fn. w.r.t.

Compute the risk fn. of the Bayes estimator
w.r.t. that prior.

Step 3: Try to find prior hyperparameters
in such a way that the risk fn. of the
Bayes estimator becomes const.

Example: $x_1, \dots, x_n \stackrel{\text{iid}}{\sim} \text{Ber}(\theta)$, $\theta \sim \text{Beta}(\alpha, \beta)$.

$$\delta_{\pi}(x) = \frac{\sum_{i=1}^n x_i + \alpha}{n + \alpha + \beta}$$

$$E_{\bar{x}|\theta} [(\delta_{\pi}(x) - \theta)^2] = E_{\bar{x}|\theta} \left[\left(\frac{\sum_{i=1}^n x_i + \alpha}{n + \alpha + \beta} - \theta \right)^2 \right]$$

$$= \text{Var}_{\bar{x}|\theta} \left(\frac{\sum_{i=1}^n x_i + \alpha}{n + \alpha + \beta} \right) + \left[E_{\bar{x}|\theta} \left(\frac{\sum_{i=1}^n x_i + \alpha}{n + \alpha + \beta} \right) - \theta \right]^2$$

$$= \frac{n\theta(1-\theta)}{(n+\alpha+\beta)^2} + \left[\frac{n\theta + \alpha}{n + \alpha + \beta} - \theta \right]^2$$

$$= \frac{n\theta(1-\theta)}{(n+\alpha+\beta)^2} + \frac{[\alpha - \theta\alpha - \theta\beta]^2}{(n + \alpha + \beta)^2}$$

$$= \frac{1}{(n+\alpha+\beta)^2} \left[\alpha^2 + \{n - 2\alpha(\alpha+\beta)\}p + \{(\alpha+\beta)^2 - n\}p^2 \right]$$

This to be a const. (not a function of p)
we must have

$$\begin{aligned} n = 2\alpha(\alpha+\beta) & \quad \text{--- --- --- } \textcircled{1} \\ \sqrt{n} = (\alpha+\beta)p & \quad \text{--- --- --- } \textcircled{2} \end{aligned}$$

Solving them we obtain $\alpha = \frac{\sqrt{n}}{2}$, $\beta = \frac{\sqrt{n}}{2}$.

Thus by our earlier thm., Beta($\frac{\sqrt{n}}{2}, \frac{\sqrt{n}}{2}$)
is ~~a~~ least favourable dist. and

$\frac{\sum_{i=1}^n x_i + \frac{\sqrt{n}}{2}}{n + \sqrt{n}}$ is a minimax estimator.

Testing of Hypothesis

statistical hypothesis testing is all about

① Beginning with a tentative idea about the unknown parameter.

② Want to test the validity of this tentative idea based on sample. Existing tentative idea, status quo: H_0 (null hypothesis), new idea: H_1 (alternative hypothesis).

③ We begin by assuming that the null hypotheses is true. Only when there is an overwhelming evidence contradicting the null do we reject it in favor of the alternative.

	H_0 is true	H_0 is false
H_0	Correct	Incorrect
Reject H_0	Incorrect	Correct

Type 1 error = $P(\text{reject } H_0 \mid H_0 \text{ is true})$

Type 2 error = $P(\text{do not reject } H_0 \mid H_0 \text{ is false})$

Type 1 error is also known as the level (size) of the test.

On the other hand, $1 - \text{Type 2 error}$ =

= $P(\text{Rejecting } H_0 \mid H_0 \text{ is false})$ = power of the test.

$$0 \leq \text{level} \leq 1, \quad 0 \leq \text{power} \leq 1$$

~~General~~ Ideally, we would like to minimize both type 1 and type 2 error. But it turns out that it is not possible to simultaneously minimize them. So, we fix level at a pre-specified value and find a test that maximizes power.

Parametric tests: $x_1, \dots, x_n \stackrel{\text{iid}}{\sim} f(x|\theta)$. We test.

$H_0: \theta \in \mathbb{H}_0$ vs. $H_1: \theta \in \mathbb{H}_1$. If \mathbb{H}_0 is singleton we will call it a simple null hypothesis. Otherwise we will call it a composite null hypothesis.

Let $R = \{\underline{x} \in \mathcal{X} \mid \text{the null hypothesis is rejected for } \underline{x}\}$.

R is called the rejection region or critical region of the test.

$\phi(\underline{x})$ = prob. of rejecting H_0 when \underline{x} is observed.
Power function of a test is given by

$$\beta(\theta) = \int \phi(\underline{x}) f(\underline{x}|\theta) d\underline{x}$$

$$= \int P(\text{rejecting } H_0 \mid \underline{x} \text{ is observed}) P(\underline{x} \text{ is observed} \mid \theta \text{ is the truth}) d\underline{x}$$

intuitively
 $= P(\text{reject } H_0 \mid \theta \text{ is the truth}).$

If level of a test is α and the power fn. of that test is given by $\beta(\theta)$ and the hypothesis is $H_0: \theta \in \mathbb{H}_0$ vs. $H_1: \theta \in \mathbb{H}_1$,
then $\beta(\theta) \leq \alpha$, $\theta \in \mathbb{H}_0$