this time: Neyman (1937) (confidence intervals) frequentist inference

AMS 206 29 Jan 19

①

next time: Fisher frequentist inference (maximum likelihood)

$$(Y_i \mid \theta\, \mathcal{B}) \stackrel{IID}{\sim} \text{Bernoulli}(\theta)$$
$$(i = 1, \dots, n)$$

$$P(Y_i = y_i \mid \theta\, \mathcal{B}) = \begin{cases} \theta & \text{if } y_i = 1 \\ 1 - \theta & \phantom{\text{if } y_i =}\ 0 \end{cases}$$

$$E(Y_i \mid \theta\, \mathcal{B}) = \sum_{\text{all } y_i} y_i\, P(Y_i = y_i \mid \theta\, \mathcal{B})$$
$$= 1 \cdot \theta + 0\,(1 - \theta)$$
$$= \theta$$

$$V(Y_i \mid \theta\, \mathcal{B}) = E(Y_i^2 \mid \theta\, \mathcal{B}) - \left[E(Y_i \mid \theta\, \mathcal{B}) \right]^2$$

②

$$E(Y_i^2 \mid \theta \mathcal{B}) = 1^2 \cdot \theta + 0^2 (1-\theta)$$

$$= \theta$$

$$V(Y_i \mid \theta \mathcal{B}) = \theta - \theta^2 = \theta(1-\theta)$$

$$SD(Y_i \mid \theta \mathcal{B}) = \sqrt{\theta(1-\theta)}$$

$$(Y_i \mid \mu \, \sigma \, \mathcal{B}) \overset{IID}{\sim} \begin{cases} E(Y_i) = \mu \\ V(Y_i) = \sigma^2 \end{cases}$$

$$(i = 1, \dots, n)$$

must assume $\sigma < \infty$

$$\bar{Y}_n = \frac{1}{n} \sum_{i=1}^{n} Y_i = \hat{\theta}$$

$$E(\bar{Y}_n) = \mu \qquad V(\bar{Y}_n) = \frac{\sigma^2}{n}$$

$$SD(\bar{Y}_n) = \frac{\sigma}{\sqrt{n}} \triangleq SE(\bar{Y}_n)$$

$$SD(\hat{\theta}) = \sqrt{\frac{\theta(1-\theta)}{n}}$$

$(Y_i \mid \theta\, B) \overset{IID}{\sim} \text{Bernoulli}(\theta)$

$i = 1, \ldots, n$

$S = \sum_{i=1}^{n} Y_i$ | $\hat{\theta} = \frac{S}{n}$

$S \sim \text{Binomial}(n, \theta)$

$\hat{\theta}$ estimate of θ

R.V.

def. $\hat{\theta}$ is unbiased for θ

iff $E_{R.S.}(\hat{\theta}) = \theta$

def. the bias of $\hat{\theta}$ is

$\text{bias}(\hat{\theta}) = E_{R.S.}(\hat{\theta}) - \theta$

$\hat{SE}$ 0.019

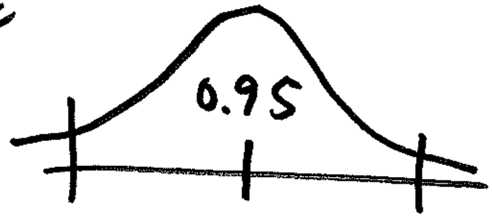

$\theta - 1.96\hat{SE}$ θ $\theta + 1.96\hat{SE}$ raw units

standard units: -1.96 0 1.96

approx. repeated-sampling dist. of $\hat{\theta}$

(PDF of r.v. $\hat{\theta}$)

SD 1

0

$$P_F\left(\boxed{\theta - 1.96\hat{SE} < \boxed{\hat{\theta}} < \theta + 1.96\hat{SE}}\right) \doteq 95\%$$

(Neyman's confidence trick) $\updownarrow$ fixed (freq)

$$P_F\Big(\underbrace{\hat{\theta} - 1.96\hat{SE}}_{0.141} < \theta < \overbrace{\hat{\theta} + 1.96\hat{SE}}^{0.216}\Big) \doteq 0.95$$

random random random

we want interval of form $\boxed{\text{some } \# < \theta < \text{some other } \#}$

$100(1-\alpha)\%$ confidence interval (CI) for θ

$$= \hat{\theta} \pm \Phi^{-1}\left(1 - \frac{\alpha}{2}\right) \cdot \sqrt{\frac{\hat{\theta}(1-\hat{\theta})}{n}}$$

(pnorm)

$\Phi(z) =$ CDF of standard normal curve

qnorm

95% CI for θ = $(0.141, 0.216)$

95% CI for θ

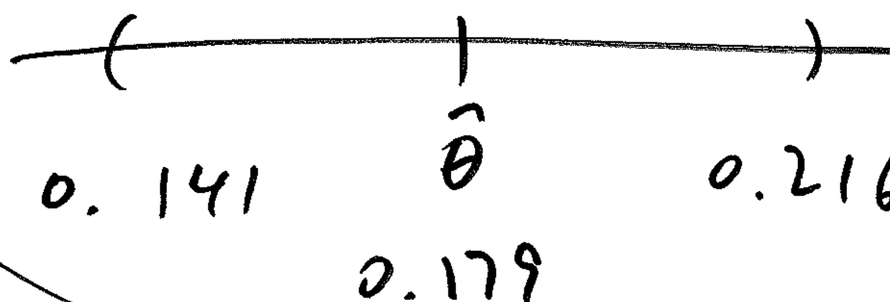

interpreting CIs

fixed unknown constant

$$P_F(0.141 < \theta < 0.216) \cancel{=} \cancel{0.95}$$

undefined

95% confidence:

about 95% of your (repeated) (imaginary) CIs would include the truth

5% of CIs are misses (error rate)

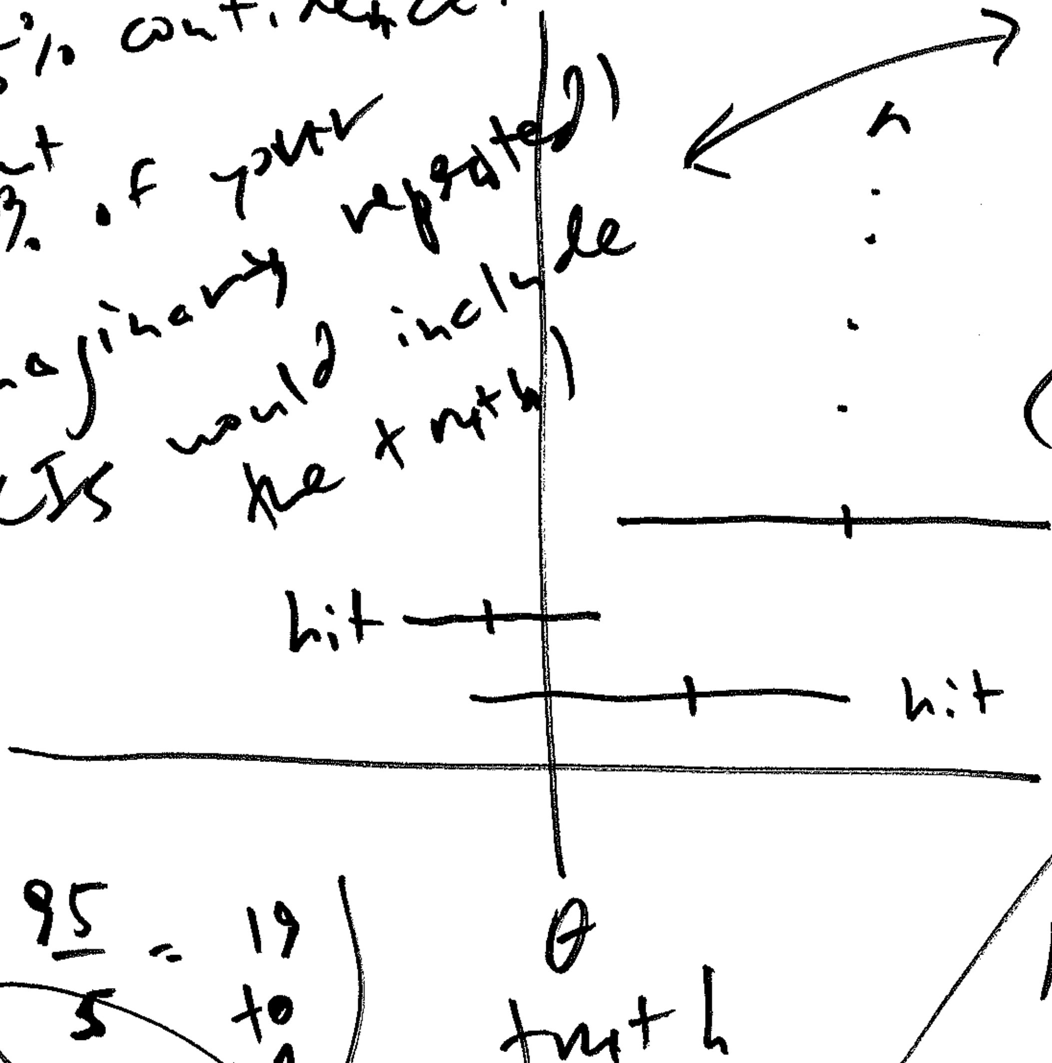

$\left(\frac{95}{5} = 19 \text{ to } 1\right)$

Mr. Neyman's confidence

is in the process of creating CIs, not in the outcome of 1 realization of that process

~~Fisher~~ 95% CI for θ

$\theta_{national}$ 0.13 | (0.141 | 0.179 | 0.216)

$\theta_{national} = 0.13$

(1) practical significance: (practsig)

$$\frac{0.179 - 0.13}{0.13} \doteq \underline{0.37};$$

DH $\hat{\theta}$ is 37% higher than $\theta_{national}$;

this is huge in practical terms

(2) statistical significance: convention

see if $\theta_{national}$ is in the 95% CI;

if not, declare $\hat{\theta}$ statistically

significantly different from $\theta_{national}$

(statsig)

$\hat{\theta}$ statsig different from $\theta_{\text{benchmark}} \leftarrow \theta_b$

$\leftrightarrow$ difference between $\hat{\theta}$ & θ_b

is hard to explain by unlucky

random sampling $\longleftrightarrow$ probably real