Administrivia

Course evaluations at 26.05% 👁👄👁 If they hit 80%, everyone gets an additional homework drop.

Recap

Thm (Markov's Inequality) For any nonnegative random variable X and $a > 0$,

$$\mathbb{P}(X \geq a) \leq \frac{\mathbb{E}[X]}{a}$$

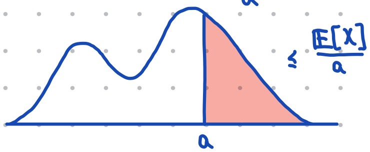

Thm (Chebyshev's Inequality) For any random variable X with expectation μ and variance $\sigma^2 < \infty$,

$$\mathbb{P}(\underbrace{|X - \mu| \geq c}_{X \text{ is more than a distance of } c \text{ from its mean}}) \leq \frac{\mathrm{Var}[X]}{c^2}$$

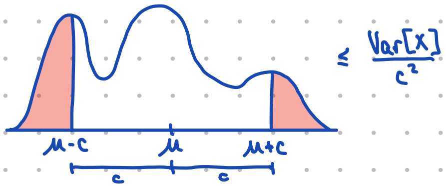

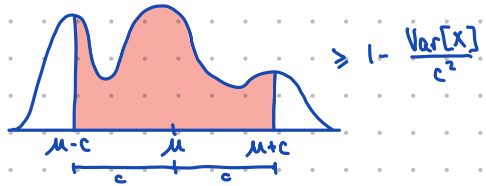

Note Normal Random Variables

For $X \sim \mathrm{Normal}(\mu, \sigma^2)$ and $Z \sim \mathrm{Normal}(0,1)$,

$$\frac{X-\mu}{\sigma} \overset{d}{=} Z \iff X \overset{d}{=} \sigma Z + \mu$$

The density of Z is symmetric $(\phi(z) = \phi(-z))$, so

$$\Phi(z) = 1 - \Phi(-z)$$

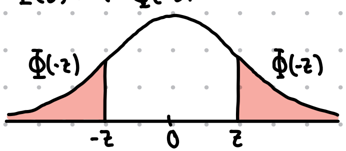

For $X_1, \dots, X_n$ iid with mean μ and variance σ^2,

$$S_n = \sum_{i=1}^{n} X_i \xrightarrow{d} \mathrm{Normal}(n\mu, n\sigma^2)$$

Estimation

Def An estimator is a random variable X that is used to estimate a fixed numerical parameter θ.
The bias of an estimator is
$$\text{Bias}[X] = \mathbb{E}[X] - \theta$$
We say an estimator is unbiased if $\text{Bias}[X] = 0 \Leftrightarrow \mathbb{E}[X] = \theta$.

Ex Suppose $X_1, \dots, X_n \overset{iid}{\sim} \text{Bernoulli}(p)$ for unknown parameter p. Construct an unbiased estimator for p.
① Note that $\mathbb{E}[X_i] = p$, so each X_i is an unbiased estimator for p.
② Consider
$$\bar{X} = \frac{1}{n}\sum_{i=1}^{n} X_i,$$
the sample mean. Then $\mathbb{E}[\bar{X}] = \mathbb{E}[X_i] = p$, so the sample mean is an unbiased estimator for p.
Which estimator is better?
① $\text{Var}[X_i] = p(1-p)$
② $\text{Var}[\bar{X}] = \frac{p(1-p)}{n} < \text{Var}[X_i]$

Ex Suppose $X_1, \dots, X_n$ are iid with unknown expectation μ and variance σ^2. Construct an unbiased estimator for μ.
Again,
$$\mathbb{E}[\bar{X}] = \mu,$$
so the sample mean is an unbiased estimator for the population mean.
$$\text{Var}[\bar{X}] = \frac{\sigma^2}{n} \to 0 \text{ as } n \to \infty$$

Note Generally, for an estimator X of θ, we want that
$\mathbb{E}[X] = \theta$
$\text{Var}[X] \to 0$ as $n \to \infty$

Chebyshev Confidence Intervals I

Def For $0 < \delta < 1$, a $(1-\delta)$ confidence interval for a fixed parameter θ is a random interval (a, b) such that

$$\mathbb{P}(a < \theta < b) = 1-\delta$$

δ is called the significance

Q We flip a biased coin that flips heads with probability p n times. Let $X_1, \dots, X_n$ be the results of the flips.

Construct an unbiased estimator for p.

By the previous example,

$$\bar{X} = \frac{1}{n} S_n$$

is an unbiased estimator for p.

Construct a $(1-\delta)$ confidence interval for p.

We construct our interval as $[\bar{X}-\varepsilon, \bar{X}+\varepsilon]$; ε is the width of our confidence interval.

$$\mathbb{P}(\bar{X}-\varepsilon \leq p \leq \bar{X}+\varepsilon) = \mathbb{P}(-\varepsilon \leq p - \bar{X} \leq \varepsilon) = \mathbb{P}(-\varepsilon \leq \bar{X}-p \leq \varepsilon) = \mathbb{P}(|\bar{X}-p| \leq \varepsilon)$$

So we want

$$\mathbb{P}(|\bar{X}-p| \leq \varepsilon) = 1-\delta$$

$$1 - \mathbb{P}(|\bar{X}-p| \leq \varepsilon) = 1-(1-\delta)$$

$$\mathbb{P}(|\bar{X}-p| > \varepsilon) = \delta$$

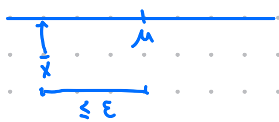

By Chebyshev,

$$\mathbb{P}(|\bar{X}-p| \geq \varepsilon) \leq \frac{Var[\bar{X}]}{\varepsilon^2}$$

So if

$$\frac{Var[\bar{X}]}{\varepsilon^2} = \delta, \text{ then } \mathbb{P}(|\bar{X}-\mu| \geq \varepsilon) \leq \delta$$

$$\frac{p(1-p)}{n\varepsilon^2} = \delta \iff \varepsilon = \sqrt{\frac{p(1-p)}{n\delta}} \leq \sqrt{\frac{1/4}{n\delta}} = \frac{1}{2\sqrt{n\delta}} \quad \text{because } p(1-p) \leq 1/4$$

Our confidence interval is

$$\bar{X} \pm \frac{1}{2\sqrt{n\delta}}$$

Chebyshev Confidence Intervals II

Q We flip a biased coin that flips heads with probability p n times. Let $X_1, \ldots, X_n$ be the results of the flips.
Suppose $n=1000$ and 120 of the flips are heads. Construct the 95% confidence interval.

$\bar{X} = 0.12$

$n = 1000$ $\Rightarrow \varepsilon = \dfrac{1}{2\sqrt{1000 \cdot 0.05}}$

$\delta = 1 - 0.95 = 0.05$

Our interval is

$$\left[0.12 - \frac{1}{2\sqrt{50}},\ 0.12 + \frac{1}{2\sqrt{50}}\right] \approx [0.049,\ 0.191]$$

Q Suppose $X_1, \ldots, X_n$ are iid with unknown expectation μ and known variance $\sigma^2 = 9$.
Find n such that a 98% confidence interval has error at most 0.01.

From before,

$$P(|\bar{X} - \mu| \geq \varepsilon) \leq \frac{Var[\bar{X}]}{\varepsilon^2}$$

We must find n such that

$$\frac{Var[\bar{X}]}{\varepsilon^2} = \frac{\sigma^2}{n\varepsilon^2} \leq \delta \iff n \geq \frac{\sigma^2}{\delta\varepsilon^2} = \frac{9}{(0.02)(0.01)^2} = 4.5 \times 10^9$$

Normal Confidence Intervals 1

Q Let $X_1, \ldots, X_n$ be iid with expectation μ and variance $\sigma^2 \in (0, \infty)$
What is the approximate distribution of $\bar{X} = \frac{1}{n}\sum X_i$ for large n?
By Central Limit Theorem, for $S_n = \sum X_i$ and $Z \sim \text{Normal}(0,1)$,

$$\frac{S_n - n\mu}{\sigma\sqrt{n}} \xrightarrow{d} Z \iff S_n \xrightarrow{d} \sigma\sqrt{n}\, Z + n\mu$$

Therefore

$$\bar{X} = \frac{1}{n} S_n \xrightarrow{d} \frac{\sigma\sqrt{n}}{n} Z + \mu \iff \bar{X} \approx \text{Normal}\left(\mu, \frac{\sigma^2}{n}\right)$$

$\bar{x}$ is approximately normal for large n.
Suppose n is large. Construct a $(1-\delta)$ confidence interval for μ, the population mean.
We construct our interval as $[\bar{x} - \varepsilon, \bar{x} + \varepsilon]$.
As before, we want

$$\mathbb{P}(|\bar{X} - \mu| \le \varepsilon) = 1 - \delta$$

However, this time we can use the distribution of $\bar{x}$.

$$\mathbb{P}\left(\left|\frac{\sigma}{\sqrt{n}} Z + \mu - \mu\right| \le \varepsilon\right) = \mathbb{P}\left(\frac{\sigma}{\sqrt{n}}|Z| \le \varepsilon\right) = \mathbb{P}\left(|Z| \le \frac{\varepsilon\sqrt{n}}{\sigma}\right) = \mathbb{P}\left(-\frac{\varepsilon\sqrt{n}}{\sigma} \le Z \le \frac{\varepsilon\sqrt{n}}{\sigma}\right) = \Phi\left(\frac{\varepsilon\sqrt{n}}{\sigma}\right) - \Phi\left(-\frac{\varepsilon\sqrt{n}}{\sigma}\right)$$

$$= \Phi\left(\frac{\varepsilon\sqrt{n}}{\sigma}\right) - \left(1 - \Phi\left(\frac{\varepsilon\sqrt{n}}{\sigma}\right)\right) = 2\Phi\left(\frac{\varepsilon\sqrt{n}}{\sigma}\right) - 1$$

This must equal $1-\delta$.

$$2\Phi\left(\frac{\varepsilon\sqrt{n}}{\sigma}\right) - 1 = 1 - \delta \iff \Phi\left(\frac{\varepsilon\sqrt{n}}{\sigma}\right) = \frac{2-\delta}{2}$$

$$\iff \frac{\varepsilon\sqrt{n}}{\sigma} = \Phi^{-1}\left(\frac{2-\delta}{2}\right)$$

$$\iff \varepsilon = \frac{\sigma}{\sqrt{n}} \Phi^{-1}\left(\frac{2-\delta}{2}\right)$$

Our confidence interval is

$$\left[\bar{X} - \frac{\sigma}{\sqrt{n}}\Phi^{-1}\left(\frac{2-\delta}{2}\right), \bar{X} + \frac{\sigma}{\sqrt{n}}\Phi^{-1}\left(\frac{2-\delta}{2}\right)\right]$$

Note When the sample size is large, the sample standard deviation is a good approximation for σ.

Normal Confidence Intervals I

Q We flip a biased coin that flips heads with probability p n times. Let $X_1, \dots, X_n$ be the results of the flips.

Suppose $n=1000$ and 120 of the flips are heads. Construct the 95% confidence interval.

$\bar{X} = 0.12$

$n = 1000$ $\Rightarrow \varepsilon = \frac{\sigma}{\sqrt{1000}} \Phi^{-1}\left(\frac{2-0.05}{2}\right) \leq \frac{1/2}{\sqrt{1000}} \Phi^{-1}\left(\frac{2-0.05}{2}\right) = 0.031$

$\delta = 0.05$

Our interval is

$$\left[0.12 - \frac{1}{2\sqrt{1000}}\Phi^{-1}(0.975),\ 0.12 + \frac{1}{2\sqrt{1000}}\Phi^{-1}(0.975)\right] \approx [0.089, 0.151]$$

Note that our interval is better than the Chebyshev interval, $[0.049, 0.191]$

Law of Large Numbers

<u>Note</u> We have seen that the variance in $\bar{X}$ decreases as the sample size increases.

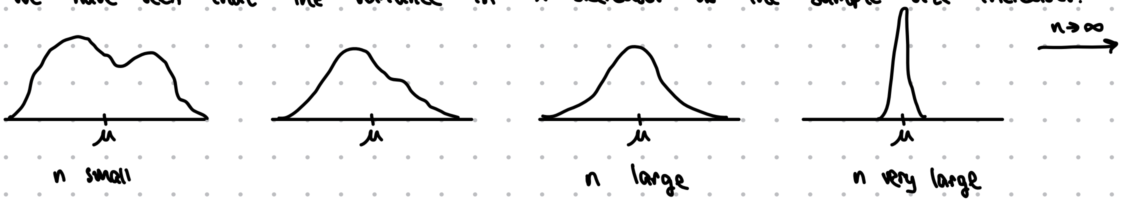

<u>Thm</u> (Law of Large Numbers) Let $X_1, \ldots, X_n$ be iid with expectation $\mu < \infty$. Let the sample mean

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

For any $\varepsilon > 0$,

$$\mathbb{P}(|\bar{X} - \mu| \leq \varepsilon) \to 1 \quad \text{as} \quad n \to \infty$$

<u>Pf</u> Let $\sigma^2 = \mathrm{Var}[X_i]$ be finite. By Chebyshev,

$$\mathbb{P}(|\bar{X} - \mu| \leq \varepsilon) \geq 1 - \frac{\mathrm{Var}[\bar{X}]}{\varepsilon^2} = 1 - \frac{\sigma^2}{n\varepsilon^2} \to 1 \quad \text{as} \quad n \to \infty$$

<u>Ex</u> I flip a biased coin with unknown probability p of heads.
Consider the distribution of $\bar{X}$ for various values of n
[Demo]