

Applications - Lecture 4

Confidence Intervals

Statistics Department

Ingénierie Mathématique & Informatique

Section	Notion	Used in
4.1	Exact/asymptotic/approximate CI, $\mathbb{P}_\theta(g(\theta) \in I_n)$	App. 1, App. 2
4.2	Pivotal functions, free random variables	App. 1 (contrast)
4.3	Asymptotic CI via a CLT-normal estimator Z_n	App. 1
4.4	Concentration inequalities (Bienayme-Chebychev, Hoeffding)	App. 2

1. **Large samples:** A/B testing CTRs with the asymptotic Wald CI.
2. **Small, bounded samples:** distribution-free confidence bounds from concentration inequalities.

App. 1 - A/B Testing in Technology Companies


Comparing two click-through rates

Problem. A company compares a new webpage (variant B) to the current one (variant A). Visitors are randomly assigned; we must *estimate* each click-through rate and *quantify uncertainty*

Model. $X_1, \dots, X_{n_A} \stackrel{\text{iid}}{\sim} \text{Bernoulli}(p_A)$, similarly $Y_1, \dots, Y_{n_B}$ for p_B . Since the Bernoulli model has **no pivotal function**, we build an **asymptotic** CI from the consistent, asymptotically normal estimator $\hat{p}_n = n^{-1} \sum_i X_i$, of asymptotic variance $V(p) = p(1 - p)$:

$$I_n = \left[\hat{p}_n - \phi_{1-\alpha/2} \sqrt{\frac{\hat{p}_n(1 - \hat{p}_n)}{n}}, \hat{p}_n + \phi_{1-\alpha/2} \sqrt{\frac{\hat{p}_n(1 - \hat{p}_n)}{n}} \right].$$

Reading the confidence intervals



Left: $\hat{p} \pm \phi_{0.975}\text{se}(\hat{p})$ for variants A ($n_A = 5\,000$) and B ($n_B = 4\,800$); true $p_A = 0.05$, $p_B = 0.065$ (dashed). Right: empirical coverage of I_n vs. n (500 Monte Carlo runs), nominal level dashed.

- Non-overlapping CIs are strong visual evidence $p_B > p_A$, but the rigorous approach builds a CI for the *difference* $p_B - p_A$.
- Coverage $\mathbb{P}_p(p \in I_n) \rightarrow 1 - \alpha$ converges **slowly**: it drops well below 95% for small n

App. 2 - Distribution-Free Bounds via Concentration Inequalities

When the CLT cannot be trusted

Problem. A small clinical trial records a bounded patient-reported outcome score $X_i \in [0, 1]$ for $n = 20$ patients. With so few, skewed observations the asymptotic Gaussian CI is unreliable, and no pivotal function is available. We instead build an **approximate** CI ($\mathbb{P}_p(p \in I_n) \geq 1 - \alpha$ for every n) from a concentration inequality.


Bienayme-Chebychev. Since $\text{Var}(X_1) \leq 1/4$ for any $X_1 \in [0, 1]$, with $r_{BC}(\alpha) = \frac{1}{2\sqrt{\alpha}}$:

$$I_n^{BC} = \left[\bar{X}_n - \frac{r_{BC}(\alpha)}{\sqrt{n}}, \bar{X}_n + \frac{r_{BC}(\alpha)}{\sqrt{n}} \right].$$

Hoeffding. With Gaussian concentration $\exp(-2r^2)$ and $r_H(\alpha) = \sqrt{-\frac{1}{2} \log(\alpha/2)}$:

$$I_n^H = \left[\bar{X}_n - \frac{r_H(\alpha)}{\sqrt{n}}, \bar{X}_n + \frac{r_H(\alpha)}{\sqrt{n}} \right].$$

Comparing the three intervals



Left: half-width vs. n at $\alpha = 0.05$ for Bienayme-Chebychev, Hoeffding, and the asymptotic Gaussian interval. Right: empirical coverage at fixed $n = 20$ ($X_i \sim \text{Beta}(2, 5)$), nominal $1 - \alpha = 0.95$ dashed.

- Hoeffding's **Gaussian concentration** $\exp(-2r^2)$ beats Chebychev's polynomial $1/r^2$: I_n^H is narrower than I_n^{BC} for every n .
- Both I_n^{BC} and I_n^H meet the $\geq 1 - \alpha$ guarantee at $n = 20$; the asymptotic interval, narrower but only asymptotically valid, **undercovers** on this skewed bounded data.