8 Confidence intervals

[CB9, BD4.4-4.5]
Let $X_1, \cdots, X_n$ be a random sample from a distribution P_θ which belongs to a family of distributions $\mathcal{P} = \{P_\theta : \theta \in \Theta\}$.

Definition 1 *For fixed $\alpha \in (0,1)$, a random interval $[T_1(X), T_2(X)]$, where $P(T_1 < T_2) = 1$, such that $P_\theta(T_1 \le \theta \le T_2) = 1 - \alpha \forall \theta \in \Theta$, is called a $100(1-\alpha)\%$ Confidence Interval (CI) for θ. Random variables T_1 and T_2 are called the lower and upper limit, respectively; $1-\alpha$ is called the confidence coefficient.*

$T_1(X)$ and $T_2(X)$ are rvs, hence their value may be different for different realizations of the sample X . For some observed samples x, the interval $[T_1(x), T_2(x)]$ may not be covering the true unknown parameter θ. If the sampling procedure is repeated a large number of times, then the proportion of samples for which the interval actually covers θ should be approximately equal to $1-\alpha$. Note that here the interval is random, while θ is an unknown constant. $1-\alpha$ is usually taken to be 0.9, 0.95 or 0.99.

8.1 Pivotal method for finding CI

Definition 2 *Let $X = (X_1, \cdots, X_n)$ be a random sample from a distribution $P_\theta, \theta \in \Theta$. A function $Q(X, \theta)$ is called a pivot for θ if its distribution is completely known.*

Example 1: Let $X_i \overset{iid}{\sim} \mathcal{N}(\mu, \sigma_0^2), \sigma_0^2$ known.
Then $Q(X, \mu) := (\bar{X} - \mu) \sim \mathcal{N}(0, \sigma_0^2/n)$ is a pivot for μ.
It is a function of the sufficient statistic $\bar{X}$.

Example 2: Let $X_i \overset{iid}{\sim} \mathcal{N}(\mu, \sigma^2)$, where both parameters are unknown.
Then $Q(X, \mu) := (\bar{X} - \mu) \sim \mathcal{N}(0, \sigma^2/n)$ is not a pivot for μ since its distribution depends on σ.
However, $Q_1(X, \mu) := \sqrt{n}(\bar{X} - \mu)/S \sim t_{n-1}$ is a pivot for μ, where $S^2 = \sum_{i=1}^n (X_i - \bar{X})^2/(n-1)$.
Also, $Q_2(X, \sigma^2) := (n-1)S^2/\sigma^2 \sim \chi^2_{n-1}$ is a pivot for σ^2.
Note, that these pivots are functions of sufficient statistics, which is usually the case.

To construct a CI we choose such values, say a and b, such that for a given $\alpha \in (0,1)$ we have $P_\theta(a \le Q(X, \theta) \le b) = 1 - \alpha \forall \theta \in \Theta$. Note that a and b are non-random since the distribution of Q is free of parameters. If Q is a strictly monotonic and continuous function of θ then this can be written as $P(T_1(X; a, b) \le \theta \le T_2(X; a, b)) = 1 - \alpha \forall \theta \in \Theta$. Then the CI for θ is $[T_1(X; a, b), T_2(X; a, b)]$.

In eg 2, Q_1 is strictly monotonic and continuous function of the parameter of interest μ. Note a and b can be chosen to be the p-th and q-th quantiles of the t distribution with $n-1$ df, such that $q - p = 1 - \alpha$. The most common choice is $q = 1 - \alpha/2$ and $p = \alpha/2$. Then a $(1-\alpha)$ level CI for μ is $\bar{X} \pm t_{n-1,\alpha/2}\frac{S}{\sqrt{n}}$.

Exercise: Show that, in the setting of example 2, using Q_2 as pivot, a $(1-\alpha)$ level confidence interval for σ^2 is

$$\left(\frac{(n-1)S^2}{\chi^2_{n-1,1-\alpha/2}}, \frac{(n-1)S^2}{\chi^2_{n-1,\alpha/2}}\right).$$

8.2 Approximate confidence intervals

If we cannot find an exact pivot, we will use an asymptotic pivot. This will often be based on the maximum likelihood estimator, which has an asymptotic normal distribution, i.e., $\sqrt{n}(\hat{\theta}-\theta) \Longrightarrow \mathcal{N}(0, 1/I(\theta))$ Here AN stands for asymptotically normal. Hence, an asymptotic pivot is a function is obtained as Q with its approximate distribution as follows:

$$Q = (X,\theta) = \sqrt{nI(\theta)}(\hat{\theta}-\theta) \sim \mathcal{N}(0,1).$$

Usually $I(\theta)$ will depend on the parameter. Then we use further approximation by substituting all the unknown parameters by their estimates (preferably consistent) to obtain $I(\hat{\theta})$. Therefore, for large n, we obtain

$$Q = (X,\theta) = \sqrt{nI(\hat{\theta})}(\hat{\theta}-\theta) \sim \mathcal{N}(0,1).$$

Example 3: Suppose that $X_i \stackrel{iid}{\sim} Poisson(\lambda)$ random variables. There is no obvious pivot in this case. The maximum likelihood estimator of λ is $\hat{\lambda} = \bar{X}$, and, for large n, we know that $\sqrt{n}(\hat{\lambda}-\lambda) \Rightarrow \mathcal{N}(0,\lambda)$. Thus, an approximate 95% confidence interval for λ is $\bar{X} \pm 1.96\sqrt{\bar{X}/n}$.

Exercise: Use the above method to find an approximate CI for binomial parameter p.

8.3 Duality of CI and Tests

There is a duality between confidence intervals and hypothesis tests.

Example 1: Let $X_1, \ldots, X_n$ be a random sample from a normal distribution having unknown mean μ and known variance σ_0^2. We are interested in testing $H : \mu = \mu_0$ versus $K : \mu \neq \mu_0$. At significance level α, consider the following test: Reject H if $\mid \bar{X} - \mu_0 \mid > \sigma_0 z_{\alpha/2}/\sqrt{n}$, and do not reject otherwise. Thus the test accepts H when

$$-\sigma_0 z_{\alpha/2}/\sqrt{n} < \bar{X} - \mu_0 < \sigma_0 z_{\alpha/2}/\sqrt{n}.$$

The latter statement is also equivalent to a $(1-\alpha)$ level CI for μ_0 given by

$$\bar{X} - \sigma_0 z_{\alpha/2}/\sqrt{n} < \mu_0 < \bar{X} + \sigma_0 z_{\alpha/2}/\sqrt{n}.$$

In other words, the CI consists precisely of all those values of μ_0 for which the null hypothesis $H : \mu = \mu_0$ is accepted. The theorems below shows that the duality between CIs and hypothesis tests holds more generally.

Theorem 1 *Suppose that for every value θ_0 in Θ there is a test at level α of the hypothesis $H : \theta = \theta_0$. Denote the acceptance region of the test by $A(\theta_0)$. Then the set $C(X) = \theta : X \in A(\theta)$ is a $(1 - \alpha)$ confidence region for θ.*

Theorem 2 *Suppose that $C(X)$ is a $(1 - \alpha)$ confidence region for θ. Then an acceptance region for a test at level α of the hypothesis $H : \theta = \theta_0$ is $A(\theta_0) = X | \theta_0 \in C(X)$.*