Confidence Intervals

An overview

- Most probability distributions are indexed by one or more *parameters*.
- For example, $N(\mu, \sigma^2)$ or $B(n, p)$.
- In significance tests, we have used *point estimators* for parameters.
- For example, for iid $Y_1, Y_2, \ldots, Y_n \sim N(\mu, \sigma^2)$, $\bar{Y}$ is a *point estimator* of μ and S^2 is a point estimator of σ^2.
- Note that $E(\bar{Y}) = \mu$ and $E(S^2) = \sigma^2$. That is, $\bar{Y}$ is an *unbiased estimator* of μ and S^2 is an unbiased estimator of σ^2.
- Another example, for $Y \sim B(n, p)$, $\hat{p} = Y/n$ is an unbiased (point) estimator of p, because $E(\hat{p}) = p$.
- Now we study *interval estimator* to give a reasonable interval for parameters (e.g. (c_1, c_2) for μ).
- The assumptions are the same as in significance testing, but we do not need a null hypothesis on the parameters (e.g. $\mu = \mu_0$).

Confidence Intervals

Normal distribution with known σ^2

- Suppose $Y_1, Y_2, \ldots, Y_n$ are iid from $N(\mu, \sigma^2)$ and σ^2 is known.
- We know that $\bar{Y}$ estimates μ, but $\bar{Y}$ can be off somewhat.
- Our goal is to get a plausible range of values for μ based on the sample data.
- Recall that $\bar{Y} \sim N(\mu, \frac{\sigma}{\sqrt{n}})$ Hence

$$\begin{aligned} 0.95 &= P(-1.96 \le Z \le 1.96) \\ &= P(-1.96 \le \frac{\bar{Y} - \mu}{\sigma/\sqrt{n}} \le 1.96) \\ &= P(-1.96\frac{\sigma}{\sqrt{n}} \le \bar{Y} - \mu \le 1.96\frac{\sigma}{\sqrt{n}}) \\ &= P(\bar{Y} - 1.96\frac{\sigma}{\sqrt{n}} \le \mu \le \bar{Y} + 1.96\frac{\sigma}{\sqrt{n}}) \end{aligned}$$

- Note that $\bar{Y}$ is random and μ is fixed.
- The interval

$$\bar{y} - 1.96\frac{\sigma}{\sqrt{n}} \le \mu \le \bar{y} + 1.96\frac{\sigma}{\sqrt{n}}$$

  is called a *95% confidence interval for* μ (or 0.95 CI for μ).

Confidence Intervals

Normal CI example

Suppose there are eight (8) observations in a sample from $N(\mu, 16)$ and the observed sample mean is $\bar{y} = 11.00$. Then $n = 8, \sigma^2 = 16$, and a 95% CI for μ is

$$11.00 - 1.96\frac{4}{\sqrt{8}} \le \mu \le 11.00 + 1.96\frac{4}{\sqrt{8}}$$

which is

$$8.23 \le \mu \le 13.77$$

or

$$11.00 \pm 2.77$$

Confidence Intervals

Remarks

- It is not true that

$$P(8.23 \le \mu \le 13.77) = 0.95$$

  because once a sample is observed, there is nothing random.
- The 95% probability has to do with the procedure. It is interpreted as, 95% of the time, the CI's calculated in this way contains μ.
- For a single case, it is interpreted as having 95% confidence that μ is between 8.23 and 13.77.
- The interval [8.23, 13.77] can be thought of as a plausible range of μ.

Confidence Intervals

Remarks

- In general, let $z_{\alpha/2}$ denote the z-score such that

$$P(-z_{\alpha/2} \le Z \le z_{\alpha/2}) = 1 - \alpha.$$

- Then we have

$$1 - \alpha = P(\bar{Y} - z_{\alpha/2}\frac{\sigma}{\sqrt{n}} \le \mu \le \bar{Y} + z_{\alpha/2}\frac{\sigma}{\sqrt{n}})$$

- A $100(1-\alpha)$*% confidence interval for* μ (or $(1-\alpha)$ CI for μ) is

$$\bar{y} - z_{\alpha/2}\frac{\sigma}{\sqrt{n}} \le \mu \le \bar{y} + z_{\alpha/2}\frac{\sigma}{\sqrt{n}}$$

  or

$$\bar{y} \pm z_{\alpha/2}\frac{\sigma}{\sqrt{n}}$$

Confidence Intervals

Normal CI example continued

Continued with the CI example that has $\bar{y} = 11.00, n = 8, \sigma^2 = 16$. Find a 90% CI for μ.

- Since $1 - \alpha = 0.90$, we have $\alpha = 0.10, \alpha/2 = 0.05, z_{\alpha/2} = 1.645$ (using Table A or the table on page 92).
- Then a 90% CI for μ is

$$11.00 - 1.645\frac{4}{\sqrt{8}} \le \mu \le 11.00 + 1.645\frac{4}{\sqrt{8}}$$

  which is

$$8.67 \le \mu \le 13.33$$

  or

$$11.00 \pm 2.33$$

- By convention, CI's are two-sided. But one-sided confidence bounds are possible.

Confidence Intervals

Normal distribution with unknown σ^2

- Suppose $Y_1, Y_2, \ldots, Y_n$ are iid from $N(\mu, \sigma^2)$ and σ^2 is unknown.
- Recall that

$$\frac{\bar{Y} - \mu}{S/\sqrt{n}} \sim T_{n-1}$$

- Let $t_{\alpha/2}$ denote the t-score such that

$$P(-t_{n-1,\alpha/2} \le T_{n-1} \le t_{n-1,\alpha/2}) = 1 - \alpha.$$

- Then we have

$$1 - \alpha = P(\bar{Y} - t_{n-1,\alpha/2}\frac{S}{\sqrt{n}} \le \mu \le \bar{Y} + t_{n-1,\alpha/2}\frac{S}{\sqrt{n}})$$

- A $(1-\alpha)$ CI for μ is

$$\bar{y} - t_{n-1,\alpha/2}\frac{s}{\sqrt{n}} \le \mu \le \bar{y} + t_{n-1,\alpha/2}\frac{s}{\sqrt{n}}$$

  or

$$\bar{y} \pm t_{n-1,\alpha/2}\frac{s}{\sqrt{n}}$$

Confidence Intervals

Tomato weight example

- Recall a random sample of $n = 16$ tomatoes that has a sample mean weight of $\bar{y} = 32.50$ gm.
- Previous we assumed that the weight of tomatoes have a normal distribution $N(\mu, (5)^2)$.
- Thus a 95% CI for μ is

$$32.50 - 1.96 \times \frac{5}{\sqrt{16}} \le \mu \le 32.50 + 1.96 \times \frac{5}{\sqrt{16}}$$

  which is [30.05, 34.95] or 32.50 ± 2.45.
- But suppose we do not know what σ is. Compute a sample variance which turns out to $s^2 = 30.02$.
- Then since $1 - \alpha = 0.95$, we have $\alpha = 0.05, \alpha/2 = 0.025$, $t_{n-1,\alpha/2} = t_{15,0.025} = 2.131$ (using Table C), and $s/\sqrt{n} = \sqrt{30.02/16} = 1.370$
- Then a 95% CI for μ is

$$32.50 - 2.131 \times 1.370 \le \mu \le 32.50 + 2.131 \times 1.370$$

  which is [29.58, 35.42] or 32.50 ± 2.92.

Confidence Intervals

Inference for unspecified distribution

- Suppose we have a large number of observations from an unspecified distribution with mean μ and variance σ^2. Also suppose that σ^2 is unknown. Our goal is to construct a CI for μ.
- By the CLT and $S^2 \approx \sigma^2$ when n is large, we have the fact that
$$\frac{\bar{Y} - \mu}{S/\sqrt{n}} \approx N(0, 1)$$
- Thus an *approximate* $(1 - \alpha)$ CI for μ is
$$\bar{y} - z_{\alpha/2}\frac{s}{\sqrt{n}} \le \mu \le \bar{y} + z_{\alpha/2}\frac{s}{\sqrt{n}}$$
- For example, for a random sample of size $n = 85$, the observed sample mean and sample variance are $\bar{y} = 25.50, s^2 = 2.83$.
- For a 95% CI, $z_{\alpha/2} = z_{0.025} = 1.96$ and
$$25.50 - 1.96 \times \sqrt{\frac{2.83}{85}} \le \mu \le 25.50 + 1.96 \times \sqrt{\frac{2.83}{85}}$$
which is $[25.14, 25.86]$ or 25.50 ± 0.36.

Confidence Intervals

Remarks

- If σ is known, then an approximate $(1 - \alpha)$ CI for μ is
$$\bar{y} - z_{\alpha/2}\frac{\sigma}{\sqrt{n}} \le \mu \le \bar{y} + z_{\alpha/2}\frac{\sigma}{\sqrt{n}}$$
- Up until this point, there is an exact relation between significance testing and CI's.
- If a $(1 - \alpha)$ CI contains the hypothesized value in $H_0 : \mu = \mu_0$, then do not reject H_0 at the α level.
- If a $(1 - \alpha)$ CI does not contain the hypothesized value in $H_0 : \mu = \mu_0$, then reject H_0 at the α level.
- Caution: This relation does not always hold (e.g. binomial).
- Suppose $Y_1, Y_2, \ldots, Y_n$ are iid from $N(\mu, \sigma^2)$. A 95% CI for σ^2 is
$$\frac{(n-1)S^2}{\chi^2_{n-1,0.025}} \le \sigma^2 \le \frac{(n-1)S^2}{\chi^2_{n-1,0.975}}$$

Confidence Intervals

Remarks

- For example, suppose a random sample of size $n = 8$ from $N(\mu, \sigma^2)$ with unknown σ^2. Suppose the observed sample and sample variance are $\bar{y} = 37.1, S^2 = 6.29$.
- For testing $H_0 : \mu = 35$ versus $H_A : \mu \neq 35$, we use t-test. The observed t-score is
$$t = \frac{\bar{y} - \mu_0}{s/\sqrt{n}} = \frac{37.1 - 35}{\sqrt{6.29/8}} = 2.368$$
and $t_{7,0.025} = 2.365$ (using Table C). Thus barely reject H_0 at the 5% level.
- A 95% CI for μ is
$$37.1 - 2.365\sqrt{6.29/8} \le \mu \le 37.1 + 2.365\sqrt{6.29/8}$$
which is $[35.003, 39.197]$, which does not contain 35 (but barely).

Confidence Intervals

Inference for binomial distribution

- Now consider $Y \sim B(n, p)$ and construct a $(1 - \alpha)$ CI for p using normal approximation.
- Recall that $\hat{p} = Y/n$ is a point estimator of p with
$$E(\hat{p}) = p, \quad Var(\hat{p}) = \frac{pq}{n}$$
- Also recall that when $np \ge 5, nq \ge 5$, we can approximate the distribution of $\hat{p}$ by
$$\hat{p}_{\text{NA}} \sim N(p, \frac{pq}{n}).$$
- For testing $H_0 : p = p_0$, use
$$\frac{\hat{p} - p_0}{\sqrt{p_0(1 - p_0)/n}} \approx N(0, 1)$$
- Now since $\hat{p} \approx p$,
$$\frac{\hat{p} - p}{\sqrt{\hat{p}(1 - \hat{p})/n}} \approx N(0, 1)$$
- Thus an approximate $(1 - \alpha)$ CI for p is
$$\hat{p} - z_{\alpha/2}\sqrt{\frac{\hat{p}(1 - \hat{p})}{n}} \le p \le \hat{p} + z_{\alpha/2}\sqrt{\frac{\hat{p}(1 - \hat{p})}{n}}$$

Confidence Intervals

Binomial CI example

In an experiment, a drug is given to treat 200 rats with a certain disease and 63 of them are cured. Let p denote the cure rate. Let Y denote the number of rats cured and assume that $Y \sim B(n, p)$. Construct a 95% CI for p.

- Here $n = 200, y = 63$, and the observed $\hat{p}$ is $y/n = 63/200 = 0.315$.
- Since $n\hat{p} = 63 > 5, n(1 - \hat{p}) = 137 > 5$, we use normal approximation.
- Since $z_{\alpha/2} = z_{0.025} = 1.96$, a 95% CI for p is
$$0.315 - 1.96 \times \sqrt{\frac{0.315 \times 0.685}{200}} \le p \le 0.315 + 1.96 \times \sqrt{\frac{0.315 \times 0.685}{200}}$$
which is $[0.251, 0.379]$ or 0.315 ± 0.064.
- Normal approximation is appropriate if $n\hat{p} \ge 5, n(1 - \hat{p}) \ge 5$.
- The relation between significance testing and CI's for p is not exact, because different variance terms are used.

Confidence Intervals

A quick summary

$Y_1, \ldots, Y_n$	n	σ known	σ unknown
$N(\mu, \sigma^2)$	small	$\bar{y} \pm z_{\alpha/2}\frac{\sigma}{\sqrt{n}}$	$\bar{y} \pm t_{n-1,\alpha/2}\frac{s}{\sqrt{n}}$
$N(\mu, \sigma^2)$	large	$\bar{y} \pm z_{\alpha/2}\frac{\sigma}{\sqrt{n}}$	$\bar{y} \pm t_{n-1,\alpha/2}\frac{s}{\sqrt{n}}$
$D(\mu, \sigma^2)$	small	no general result	no general result
$D(\mu, \sigma^2)$	large	$\bar{y} \pm z_{\alpha/2}\frac{\sigma}{\sqrt{n}}$ by CLT	$\bar{y} \pm z_{\alpha/2}\frac{s}{\sqrt{n}}$ by CLT

Confidence Intervals

Key R commands

```
> # Normal CI example
> ybar = 11
> sd = 4
> n = 8
>
> alpha = 0.05
> z = qnorm(alpha/2, lower.tail=F)
> z*sd/sqrt(n)
[1] 2.771808
> c(ybar-z*sd/sqrt(n),ybar+z*sd/sqrt(n))
[1]  8.228192 13.771808
> # Normal CI example continued
> alpha = 0.10
> z = qnorm(alpha/2, lower.tail=F)
> z*sd/sqrt(n)
[1] 2.326174
> c(ybar-z*sd/sqrt(n),ybar+z*sd/sqrt(n))
[1]  8.673826 13.326174
>
```


Confidence Intervals

Key R commands

```
> # Tomato weight example
> ybar = 32.5
> n = 16
> alpha = 0.05
>
> sd = 5
> z = qnorm(alpha/2, lower.tail=F)
> z*sd/sqrt(n)
[1] 2.449955
> c(ybar-z*sd/sqrt(n),ybar+z*sd/sqrt(n))
[1] 30.05005 34.94995
> # Fruit can example continued
> sd = sqrt(30.02)
> sd/sqrt(n)
[1] 1.369763
> t = qt(alpha/2, n-1, lower.tail=F)
> t
[1] 2.131450
> t*sd/sqrt(n)
[1] 2.91958
> c(ybar-t*sd/sqrt(n),ybar+t*sd/sqrt(n))
[1] 29.58042 35.41958
> # or directly t.test(x, conf.level = 0.95)
>
> #CLT CI example
> ybar = 25.5
> sd = sqrt(2.83)
> n = 85
> alpha = 0.05
> z = qnorm(alpha/2, lower.tail=F)
> z*sd/sqrt(n)
[1] 0.3576283
> c(ybar-z*sd/sqrt(n),ybar+z*sd/sqrt(n))
[1] 25.14237 25.85763
>
```

Confidence Intervals

Key R commands

```
> #Binomial CI example
> y=63
> n=200
> phat = y/n
> alpha = 0.05
> z = qnorm(alpha/2, lower.tail=F)
> z*sqrt(phat*(1-phat)/n)
[1] 0.06437743
> c(phat-z*sqrt(phat*(1-phat)/n),phat+z*sqrt(phat*(1-phat)/n))
[1] 0.2506226 0.3793774
> #or, directly
> prop.test(y, n, conf.level=0.95)

        1-sample proportions test with continuity correction

data:  y out of n, null probability 0.5
X-squared = 26.645, df = 1, p-value = 2.445e-07
alternative hypothesis: true p is not equal to 0.5
95 percent confidence interval:
 0.2523053 0.3849353
sample estimates:
    p
0.315
```