Inference for Single Samples

Corresponds to
Chapter 7 of
Tamhane and Dunlop

Slides prepared by Elizabeth Newton (MIT),
with some slides by Ramón V. León (University of Tennessee)

Inference About the Mean and Variance of a Normal Population

Applications:

- Monitor the mean of a manufacturing process to determine if the process is under control
- Evaluate the precision of a laboratory instrument measured by the variance of its readings
- Prediction intervals and tolerance intervals which are methods for estimating future observations from a population.

By using the central limit theorem (CLT), inference procedures for the mean of a normal population can be extended to the mean of a non-normal population when a large sample is available

Inferences on Mean (Large Samples)

- Inferences on μ will be based on the sample mean $\bar{X}$, which is an unbiased estimator of μ with variance $\dfrac{\sigma^2}{n}$.
- For large sample size *n*, the CLT tells us that $\bar{X}$ is approximately $N\left(\mu, \sigma^2/n\right)$ distributed, even if the population is not normal.
- Also for large *n*, the sample variance s^2 may be taken as an accurate estimator of σ^2 with neglible sampling error. If $n \geq 30,$ we may assume that $\sigma \simeq s$ in the formulas.

Pivots

- Definition: Casella & Berger, p. 413

- E.g. $Z = \dfrac{\overline{X} - \mu}{\sigma / \sqrt{n}} \sim N(0,1)$
- Allow us to construct confidence intervals on parameters.

Confidence Intervals on the Mean: Large Samples

$$P\left[-z_{\alpha/2} \le Z = \frac{\bar{X}-\mu}{\sigma/\sqrt{n}} \le z_{\alpha/2}\right] = 1-\alpha$$

Note: $z_{\alpha/2}$ = -qnorm(α/2)

(See Figure 2.15 on page 56 of the course textbook.)

Confidence Intervals on the Mean

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

$$\bar{x} - z_{\alpha}\frac{\sigma}{\sqrt{n}} \leq \mu \text{ (Lower One-Sided CI)}$$

$$\mu \leq \bar{x} + z_{\alpha}\frac{\sigma}{\sqrt{n}} \quad \text{(Upper One-Sided CI)}$$

$\frac{\sigma}{\sqrt{n}}$ is the standard error of the mean

Confidence Intervals in S-Plus

```
t.test(lottery.payoff)

     One-sample t-Test

data:  lottery.payoff
t = 35.9035, df = 253, p-value = 0
alternative hypothesis: true mean is not equal to 0
95 percent confidence interval:
 274.4315 306.2850
sample estimates:
 mean of x
  290.3583
```

Sample Size Determination for a z-interval

- Suppose that we require a (1-α)-level two-sided CI for μ of the form $\left[\bar{x} - E,\ \bar{x} + E\ \right]$ with a margin of error E
- Set $E = z_{\alpha/2}\dfrac{\sigma}{\sqrt{n}}$ and solve for n, obtaining $n = \left[\dfrac{z_{\alpha/2}\sigma}{E}\right]^2$

- Calculation is done at the design stage so a sample estimate of σ is not available.
- An estimate for σ can be obtained by anticipating the range of the observations and dividing by 4.

  Based on assuming normality since then 95% of the observation are expected to fall in $\left[\mu - 2\sigma, \mu + 2\sigma\right]$

Example 7.1 (Airline Revenue)

See Example 7.1, “Airline Revenue,” on page 239 of the course textbook.

Example 7.2 – Strength of Steel Beams

See Example 7.2 on page 240 of the course textbook.

Power Calculation for One-sided Z-tests

$$\pi(\mu) = \mathrm{P}[\text{Test rejects } H_0 \mid \mu]$$

Testing $H_o : \mu \le \mu_o$ vs. $H_1 : \mu > \mu_0$

For the power function of the α-level upper one sided z-test derivation, see Equation 7.7 in the course textbook.

Illustration of calculation on next page

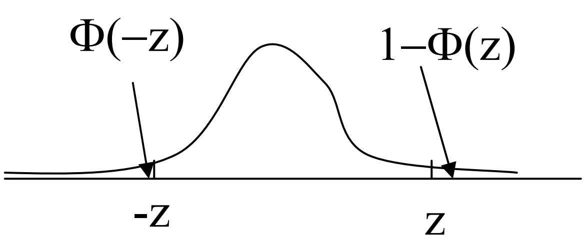

Power Calculation for One-sided Z-tests

(See Figure 7.1 on page 243 of the course textbook.)

p.d.f. curves of

$$\bar{X} \sim N\left(\mu, \frac{\sigma^2}{n}\right)$$

Power Functions Curves

See Figure 7.2 on page 243 of the course textbook.

Notice how it is easier to detect a big difference from μ_0.

Example 7.3 (SAT Couching: Power Calculation)

See Example 7.3 on page 244 of the course textbook.

$$\pi(\mu) = \Phi\left[-z_{\alpha} + \frac{(\mu - \mu_0)\sqrt{n}}{\sigma} \right]$$

Power Calculation Two-Sided Test

(See Figure 7.3 on page 245 of the course textbook.)

Power Curve for Two-sided Test

(See Figure 7.4 on page 246 of the course textbook.)

It is easier to detect large differences from the null hypothesis

Larger samples lead to more powerful tests

Power as a function of μ and n, μ_0=0, σ=1
Uses function persp in S-Plus

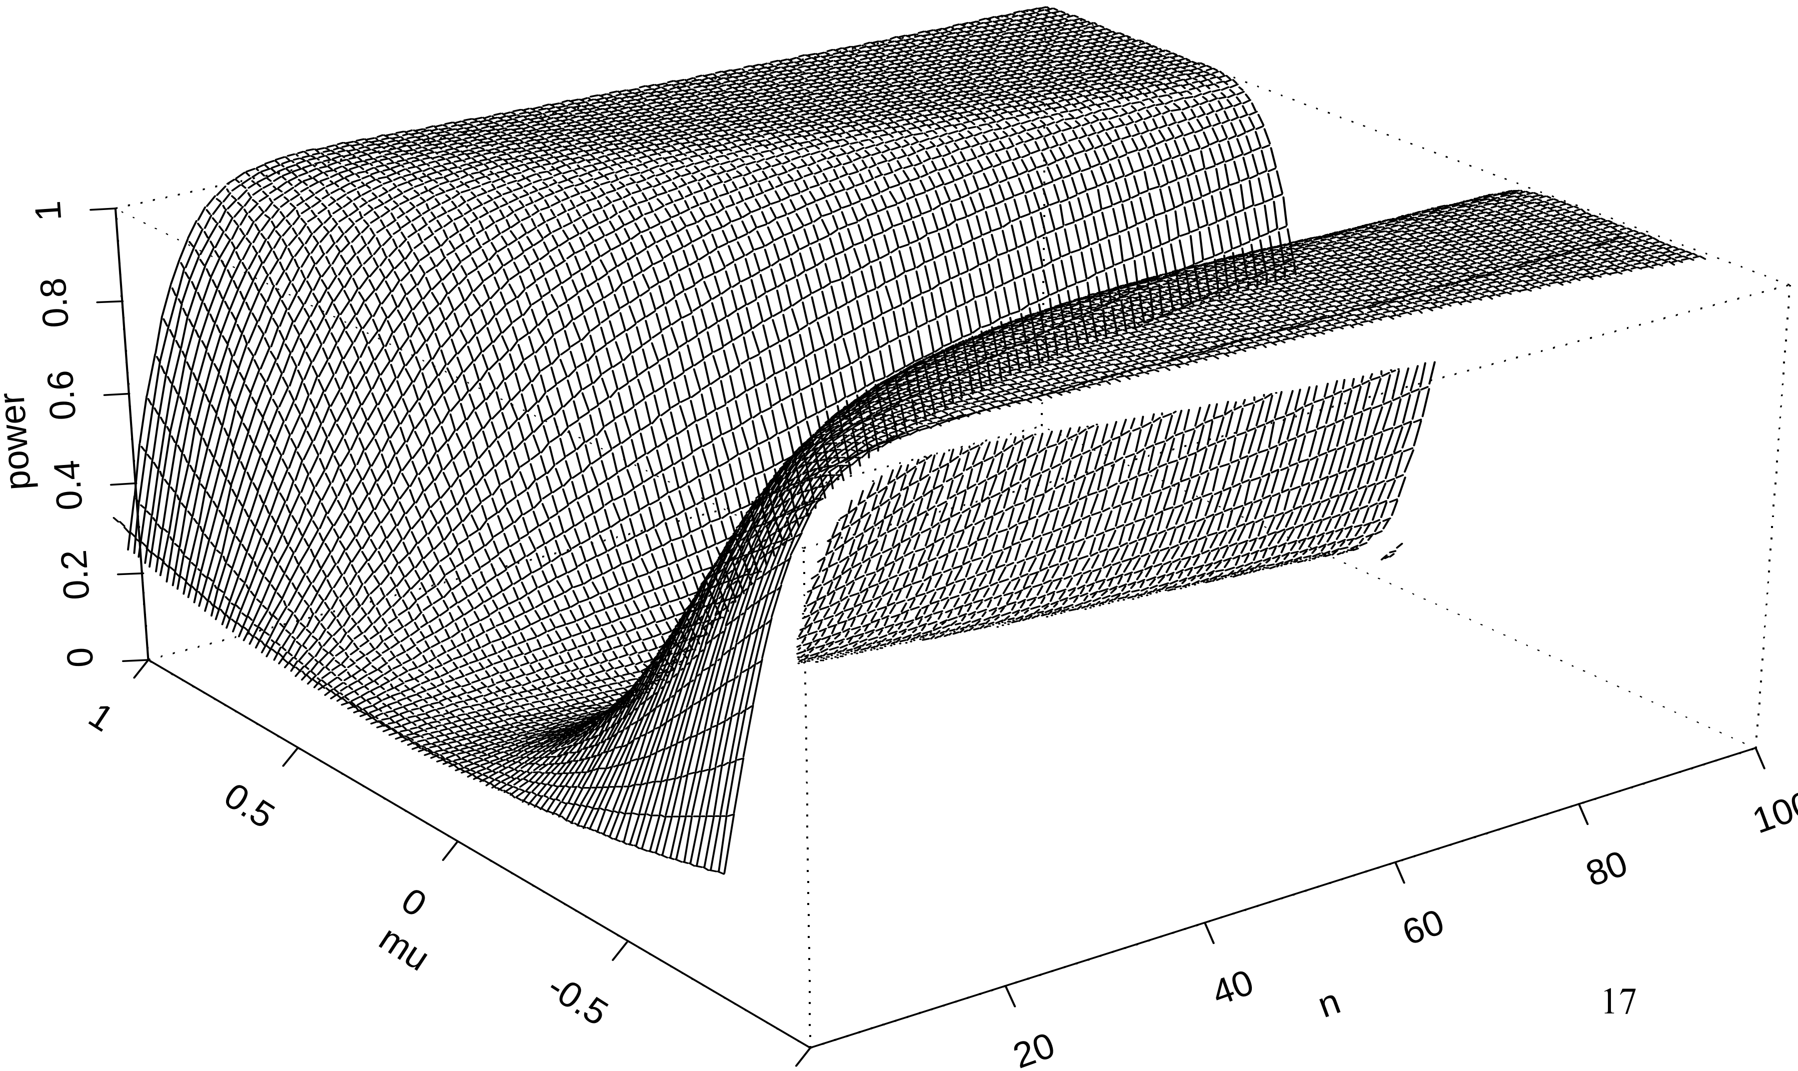

Sample Size Determination for a One-Sided *z*-Test

- Determine the sample size so that a study will have sufficient power to detect an effect of practically important magnitude
- If the goal of the study is to show that the mean response μ under a treatment is higher than the mean response μ_0 without the treatment, then $\mu - \mu_0$ is called the **treatment effect**
- Let $\delta > 0$ denote a practically important treatment effect and let $1-\beta$ denote the minimum power required to detect it. The goal is to find the minimum sample size n which would guarantee that an α-level test of H_0 has at least $1-\beta$ power to reject H_0 when the treatment effect is at least δ.

Sample Size Determination for a One-sided Z-test

Because Power is an increasing function of $\mu - \mu_0$, it is only necessary to find *n* that makes the power $1-\beta$ at $\mu = \mu_0 + \delta$.

$$\pi(\mu_0 + \delta) = \Phi\left(-z_\alpha + \frac{\delta\sqrt{n}}{\sigma}\right) = 1 - \beta \text{ [See Equation (7.7), Slide 11]}$$

Since $\Phi(z_\beta) = 1 - \beta$ we have $-z_\alpha + \dfrac{\delta\sqrt{n}}{\sigma} = z_\beta$.

Solving for n, we obtain

$$n = \left[\frac{(z_\alpha + z_\beta)\sigma}{\delta}\right]^2$$

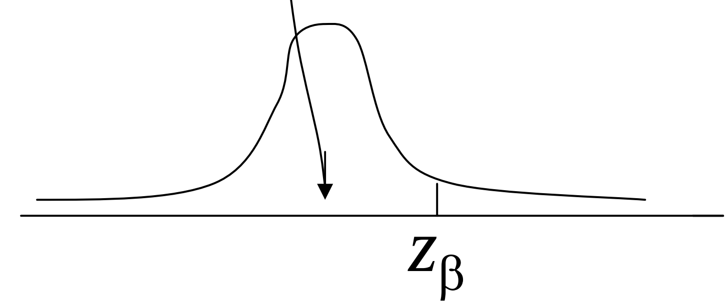

Example 7.5 (SAT Coaching: Sample Size Determination

See Example 7.5 on page 248 of the course textbook.

Sample Size Determination for a Two-Sided z-Test

$$n \doteq \left[\frac{\left(z_{\alpha/2} + z_{\beta} \right) \sigma}{\delta} \right]^2$$

Read on you own the derivation on pages 248-249

See Example 7.6 on page 249 of the course textbook.

Read on your own Example 7.4 (page 246)

Power and Sample Size in S-Plus

```
normal.sample.size(mean.alt = 0.3)
  mean.null sd1 mean.alt delta alpha power n1
          0   1      0.3   0.3  0.05   0.8 88

> normal.sample.size(mean.alt = 0.3,n1=100)
  mean.null sd1 mean.alt delta alpha  power  n1
          0   1      0.3   0.3  0.05 0.8508 100
```

Inference on Mean (Small Samples)

The sampling variability of s^2 may be sizable if the sample is small (less than 30). Inference methods must take this variability into account when σ^2 is unknown .

Assume that $X_1,..., X_n$ is a random sample from an $N(\mu, \sigma^2)$ ditribution. Then $T = \dfrac{\bar{X} - \mu}{S/\sqrt{n}}$ has a t-distribution with n-1 degrees of freedom (d.f.)

(Note that T is a pivot)

Confidence Intervals on Mean

$$1-\alpha = P\left[-t_{n-1,\alpha/2} \le T = \frac{\bar{X}-\mu}{S/\sqrt{n}} \le t_{n-1,\alpha/2}\right]$$

$$= P\left[\bar{X} - t_{n-1,\alpha/2}\frac{S}{\sqrt{n}} \le \mu \le \bar{X} + t_{n-1,\alpha/2}\frac{S}{\sqrt{n}}\right]$$

$$\bar{X} - t_{n-1,\alpha/2}\frac{S}{\sqrt{n}} \le \mu \le \bar{X} + t_{n-1,\alpha/2}\frac{S}{\sqrt{n}}$$ [Two-Sided 100(1-α)% CI]

$t_{n-1,\alpha/2} > z_{\alpha/2} \Rightarrow t-$interval is wider on the average than z-interval

Example 7.7, 7.8, and 7.9

See Examples 7.7, 7.8, and 7.9 from the course textbook.

Inference on Variance

Assume that $X_1, ..., X_n$ is a random sample from an $N(\mu, \sigma^2)$ distribution

$$\chi^2 = \frac{(n-1)S^2}{\sigma^2}$$ has a Chi-square distribution with n - 1 d.f.

(See Figure 7.8 on page 255 of
the course textbook)

$$1-\alpha = P\left[\chi^2_{n-1,1-\frac{\alpha}{2}} \le \frac{\left(n-1\right)S^2}{\sigma^2} \le \chi^2_{n-1,\frac{\alpha}{2}}\right]$$

CI for σ^2 and σ

The $100(1-\alpha)\%$ two-sided CI for σ^2 (Equation 7.17 in course textbook):

$$\frac{(n-1)s^2}{\chi^2_{n-1,\frac{\alpha}{2}}} \le \sigma^2 \le \frac{(n-1)s^2}{\chi^2_{n-1,1-\frac{\alpha}{2}}}$$

The $100(1-\alpha)\%$ two-sided CI for σ (Equation 7.18 in course textbook):

$$s\sqrt{\frac{n-1}{\chi^2_{n-1,\frac{\alpha}{2}}}} \le \sigma \le s\sqrt{\frac{n-1}{\chi^2_{n-1,1-\frac{\alpha}{2}}}}$$

Hypothesis Test on Variance

See Equation 7.21 on page 256 of the course textbook for an explanation of the chi-square statistic:

$$\chi^2 = \frac{(n-1)s^2}{\sigma_0^2}$$

Prediction Intervals

- Many practical applications call for an interval estimate of
  - an individual (future) observation sampled from a population
  - rather than of the mean of the population.
- An interval estimate for an individual observation is called a **prediction interval**

<u>Prediction Interval Formula:</u>

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

Confidence vs. Prediction Interval

Prediction interval of a single future observation:

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

As $n \to \infty$ interval converges to $[\mu - z_{\alpha/2}\sigma, \mu + z_{\alpha/2}\sigma]$

Confidence interval for μ:

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

As $n \to \infty$ interval converges to single point μ

Example 7.12: Tear Strength of Rubber

See Example 7.12 on page 259 of the course textbook.

Run chart shows process is predictable.

Tolerance Intervals

Suppose we want an interval which will contain at least .90 = 1-γ of the strengths of the future batches (observations) with 95% = 1-α confidence

Using Table A.12 in the course textbook:
1-α = 0.95
1-γ = 0.90
n = 14
So, the critical value we want is 2.529.

$$[\bar{x} - Ks, \bar{x} + Ks] = 33.712 \pm 2.529 \times 0.798 = [31.694, 35.730]$$

Note that this statistical interval is even wider than the prediction interval