

STAT 516 Lec 01

Inference on the mean and variance of a Normal population

Karl Gregory

2026-08-26

Setup


Throughout let $X_1, \dots, X_n \stackrel{\text{ind}}{\sim} \mathcal{N}(\mu, \sigma^2)$.

In this lecture we review how to:

1. Estimate μ and σ^2 .
2. Build confidence intervals for μ and σ^2 .
3. Test hypotheses concerning μ and σ^2 .
4. Choose the sample size.

We call $X_1, \dots, X_n$ a random sample.

Golden ratio example:




A class of $n = 27$ students measured B/A on themselves:

```
gr <- c(1.66, 1.61, 1.62, 1.69, 1.58, 1.43, 1.66,  
        1.69, 1.58, 1.20, 1.52, 1.60, 1.55, 1.67,  
        1.77, 1.50, 1.64, 1.54, 1.40, 1.36, 1.50,  
        1.40, 1.35, 1.48, 1.64, 1.91, 1.70)
```

What is the true mean of B/A ? Could it be the golden ratio?!?

Check if B/A measurements come from a Normal distribution.

```
qqnorm(scale(gr))  
abline(0,1)
```



Estimation

Based on $X_1, \dots, X_n$, define the sample statistics


$$\begin{aligned} \blacktriangleright \bar{X}_n &= \frac{1}{n} \sum_{i=1}^n X_i \\ \blacktriangleright S_n^2 &= \frac{1}{n-1} \sum_{i=1}^n (X_i - \bar{X}_n)^2 \end{aligned}$$

Then $\bar{X}_n$ and S_n^2 are unbiased estimators of μ and σ^2 , respectively.

Golden ratio example (cont):

We have $\bar{X}_n = \text{mean}(\text{gr}) = 1.565$ and $S_n^2 = \text{var}(\text{gr}) = 0.0219$.

```
hist(gr)
```



Important sampling distribution results

Provided $X_1, \dots, X_n \stackrel{\text{ind}}{\sim} \mathcal{N}(\mu, \sigma^2)$, we have

$$\blacktriangleright \frac{\bar{X}_n - \mu}{\sigma/\sqrt{n}} \sim \mathcal{N}(0, 1)$$

$$\blacktriangleright \frac{(n-1)S_n^2}{\sigma^2} \sim \chi_{n-1}^2$$

$$\blacktriangleright \frac{\bar{X}_n - \mu}{S_n/\sqrt{n}} \sim t_{n-1}$$

Discuss: Anatomy of chi-square and t random variables

$$\text{▶ } Z_1, \dots, Z_m \stackrel{\text{ind}}{\sim} \mathcal{N}(0, 1) \implies Z_1^2 + \dots + Z_m^2 \sim \chi_m^2.$$

$$\text{▶ } Z \sim \text{Normal}(0, 1) \perp\!\!\!\perp W \sim \chi_m^2 \implies \frac{Z}{\sqrt{W/m}} \sim t_m.$$

Relate these to the results on the previous slide.


Simulation illustrating sampling distribution results:

```
sims <- 1000
mu <- 1
sigma <- 1/2
n <- 8
Tn <- numeric(sims)
Wn <- numeric(sims)
for(s in 1:sims){


  X <- rnorm(n,mu,sigma)
  sn <- sd(X)
  xbar <- mean(X)
  Tn[s] <- sqrt(n)*(xbar - mu) / sn
  Wn[s] <- (n-1)*sn^2 / sigma^2

}
```

```
hist(Tn,freq = FALSE,breaks = 50)  
x <- seq(-4,4,length = 500)  
lines(dt(x,n-1)~x)
```



```
hist(Wn,freq = FALSE,breaks = 50)
x <- seq(0,max(Wn),length = 500)
lines(dchisq(x,n-1)~x)
```



Confidence intervals for the mean and variance

The sampling distribution results give $(1 - \alpha)100\%$ CIs as

- ▶ $\bar{X}_n \pm t_{n-1, \alpha/2} \frac{S_n}{\sqrt{n}}$ for μ .
- ▶ $\left(\frac{(n-1)S_n^2}{\chi_{n-1, \alpha/2}^2}, \frac{(n-1)S_n^2}{\chi_{n-1, 1-\alpha/2}^2} \right)$ for σ^2 .

Exercise: Derive the above.

Golden ratio example (cont):

Build 95% CIs for population mean and variance of B/A values:

```
alpha <- 0.05
n <- length(gr)

lomu <- mean(gr) - qt(1-alpha/2,n-1) * sd(gr)/sqrt(n)
upmu <- mean(gr) + qt(1-alpha/2,n-1) * sd(gr)/sqrt(n)

losgs <- (n-1) * var(gr) / qchisq(1-alpha/2,n-1)
upsgs <- (n-1) * var(gr) / qchisq(alpha/2,n-1)
```

The 95% CI for μ is (1.506,1.623). For σ^2 it is (0.014,0.041).

Testing hypotheses about the mean

Consider testing hypotheses about μ of the form

$$\begin{array}{lll} H_0: \mu \geq \mu_0 & \text{or} & H_0: \mu = \mu_0 & \text{or} & H_0: \mu \leq \mu_0 \\ H_1: \mu < \mu_0 & & H_1: \mu \neq \mu_0 & & H_1: \mu > \mu_0. \end{array}$$

Reject or fail to reject H_0 based on the value of the test statistic

$$T_{\text{stat}} = \frac{\bar{X}_n - \mu_0}{S_n / \sqrt{n}}.$$

Rejection rules for the above at significance level α are

$$T_{\text{stat}} < -t_{n-1, \alpha} \quad \text{or} \quad |T_{\text{stat}}| > t_{n-1, \alpha/2} \quad \text{or} \quad T_{\text{stat}} > t_{n-1, \alpha}.$$

The corresponding p-values are, with $T \sim t_{n-1}$, the probabilities

$$P(T < T_{\text{stat}}) \quad \text{or} \quad 2 \times P(T > |T_{\text{stat}}|) \quad \text{or} \quad P(T > T_{\text{stat}}).$$

Golden ratio example (cont):

Test $H_0: \mu = 1.618$ vs $H_1: \mu \neq 1.618$ at $\alpha = 0.05$ based on data.

```
alpha <- 0.05  
Tstat <- (mean(gr) - 1.618) / (sd(gr) / sqrt(n))  
abs(Tstat) > qt(1-alpha/2,n-1)
```

```
[1] FALSE
```

Fail to reject H_0 since $T_{\text{stat}} = -1.866$ is smaller in absolute value than $t_{n-1,\alpha/2} = 2.056$.

```
pval <- 2*(1 - pt(abs(Tstat),n-1))
```

Equivalently, the p-value, which is 0.073, is greater than $\alpha = 0.05$.

The `t.test()` function in R

The function `t.test()` tests $H_0: \mu = 0$ vs $H_1: \mu \neq 0$ by default.

```
t.test(gr)
```

One Sample t-test

```
data: gr
t = 54.902, df = 26, p-value < 2.2e-16
alternative hypothesis: true mean is not equal to 0
95 percent confidence interval:
 1.506228 1.623401
sample estimates:
mean of x
 1.564815
```

The `t.test()` function in R

Now test $H_0: \mu = 1.618$ versus $H_1: \mu \neq 1.618$, ask for 99% CI.

```
t.test(gr, mu = 1.618, conf.level = 0.99)
```

One Sample t-test

```
data: gr
t = -1.866, df = 26, p-value = 0.07336
alternative hypothesis: true mean is not equal to 1.618
99 percent confidence interval:
 1.485616 1.644013
sample estimates:
mean of x
 1.564815
```

The `t.test()` function in R

Now test $H_0: \mu \leq 1.618$ versus $H_1: \mu > 1.618$.

```
t.test(gr, mu = 1.618, alternative = "greater")
```

One Sample t-test

```
data: gr
t = -1.866, df = 26, p-value = 0.9633
alternative hypothesis: true mean is greater than 1.618
95 percent confidence interval:
 1.516202      Inf
sample estimates:
mean of x
 1.564815
```

Testing hypotheses about the variance

Consider testing hypotheses about σ^2 of the form

$$\begin{array}{ll} H_0: \sigma^2 \geq \sigma_0^2 & \text{or} \quad H_0: \sigma^2 \leq \sigma_0^2 \\ H_1: \sigma^2 < \sigma_0^2 & \quad \quad H_1: \sigma^2 > \sigma_0^2 \end{array}$$

Reject or fail to reject H_0 based on the value of the test statistic

$$W_{\text{stat}} = \frac{(n-1)S_n^2}{\sigma_0^2}.$$

Rejection rules for the above at significance level α are

$$W_{\text{stat}} < \chi_{n-1, 1-\alpha}^2 \quad \text{or} \quad W_{\text{stat}} > \chi_{n-1, \alpha}^2$$

The corresponding p-values are, with $W \sim \chi_{n-1}^2$, the probabilities

$$P(W < W_{\text{stat}}) \quad \text{or} \quad P(W > W_{\text{stat}}).$$

Golden ratio example (cont):

Test $H_0: \sigma^2 \geq 0.03$ vs $H_1: \sigma^2 < 0.03$ at $\alpha = 0.05$ based on data.

```
alpha <- 0.05  
Wstat <- (n-1)*var(gr) / 0.03  
Wstat < qchisq(alpha,n-1)
```

[1] FALSE

FTR H_0 since $W_{\text{stat}} = 19.009$ is not less than $\chi_{n-1,1-\alpha}^2 = 15.379$.

```
pval <- pchisq(Wstat,n-1)
```

Equivalently, the p-value, which is 0.164, is greater than $\alpha = 0.05$.

Pallet weights example

The weights in lbs of several wooden pallets were recorded.



```
wt <- c(41,34,40,44,33,42,52,38,32,31,31,35,  
        39,44,42,42,35,33,40,48,51,32,41,35,  
        38,48,37,35,42,41,40,47,40,46,33,38,  
        51,39,40)
```

Pallet weights example continued

- ▶ Check if we can assume normally distributed pallet weights.
- ▶ Build a 95% CI for the mean pallet weight assuming $\sigma = 5$.
- ▶ Build a 95% CI for the mean pallet weight assuming σ is unknown.
- ▶ Test at $\alpha = 0.05$ whether the mean pallet weight is equal to 48 lbs.
- ▶ Test at $\alpha = 0.05$ whether the mean pallet weight is less than 40 lbs.

Sample size calculations

We can choose a sample size based on the desired:

- a. Width of a confidence interval.
- b. Power of a test to reject H_0 when it is false.

Sample size required to achieve desired CI width

A CI for μ takes the form $\bar{X}_n \pm M$, where

- ▶ $M = z_{\alpha/2} \frac{\sigma}{\sqrt{n}}$ if σ is known
- ▶ $M = t_{n-1, \alpha/2} \frac{S_n}{\sqrt{n}}$ if σ is unknown

For ease, use the “ σ -known” version.

If one wants $M \leq M^*$, find smallest n such that $z_{\alpha/2} \frac{\sigma}{\sqrt{n}} \leq M^*$.

So take $n = \left\lceil \left(\frac{z_{\alpha/2} \sigma}{M^*} \right)^2 \right\rceil$, where $\lceil \cdot \rceil$ rounds up.

Must put in a guess for σ .

Golden ratio example (cont):

Find n required to make the 95% CI for μ no wider than 0.08.

```
alpha <- 0.05
M <- 0.08/2
sigma_guess <- sd(gr)
nr <- ceiling((qnorm(1-alpha/2) * sigma_guess / M)^2)
nr
```

```
[1] 53
```

Sample size required to achieve desired power

The power of a test is the probability with which it rejects H_0 .

For tests of H_0 concerning the mean μ we write the power as

$$\gamma(\mu) = P(\text{Reject } H_0 \text{ when true mean is } \mu) = P_\mu(\text{Reject } H_0).$$

So the power depends on the true value of μ , i.e. is a function of μ .

Exercise: Derive the power functions for the tests of

$$\begin{array}{lll} H_0: \mu \geq \mu_0 & \text{and} & H_0: \mu = \mu_0 & \text{and} & H_0: \mu \leq \mu_0 \\ H_1: \mu < \mu_0 & & H_1: \mu \neq \mu_0 & & H_1: \mu > \mu_0 \end{array}$$

with the rejection rules

$$Z_{\text{stat}} < -z_\alpha \quad \text{and} \quad |Z_{\text{stat}}| > z_{\alpha/2} \quad \text{and} \quad Z_{\text{stat}} > z_\alpha,$$


respectively, where $Z_{\text{stat}} = \frac{\bar{X}_n - \mu_0}{\sigma/\sqrt{n}}$ (σ -known case).

Plot of power curves for right-, left-, and two-sided tests

Set $\mu_0 = 1.618$, $n = 27$, and $\alpha = 0.05$ and use $\sigma = 0.148$.

```
alpha <- 0.05
mu0 <- 1.618
sigma <- sd(gr)
n <- length(gr)
mu <- seq(mu0-4*sigma/sqrt(n),mu0+4*sigma/sqrt(n),length=500)
za <- qnorm(1-alpha)
za2 <- qnorm(1-alpha/2)
d <- sqrt(n) * (mu - mu0) / sigma
rp <- 1 - pnorm(za - d)
lp <- pnorm(-za - d)
rp2 <- 1 - pnorm(za2 - d)
lp2 <- pnorm(-za2 - d)
tsp <- lp2 + rp2
```

```
plot(rp ~ mu, type = "l", ylab = "power", xlab = "mu")  
lines(lp ~ mu, lty = 2)  
lines(tsp ~ mu, lty = 3)
```



Power curve for right-sided test at various sample sizes

Test $H_0: \mu \leq 1.618$ vs $H_1: \mu > 1.618$. Use $\alpha = 0.05$ and $\sigma = 0.148$.


```
alpha <- 0.05
mu0 <- 1.618
sigma <- sd(gr)
n <- length(gr)
mu <- seq(mu0-1*sigma/sqrt(n),mu0+4*sigma/sqrt(n),length=500)
za <- qnorm(1-alpha)

# various sample sizes
nn <- c(10,20,30,40,50,60)
rp <- matrix(NA,500,length(nn))
for(j in 1:length(nn)){
  d <- sqrt(nn[j]) * (mu - mu0) / sigma
  rp[,j] <- 1 - pnorm(za - d)
}
```

```

plot(NA,xlim = range(mu), ylim = c(0,1), ylab = "power", xlab = "mu")
for(j in 1:length(nn)) lines(rp[,j] ~ mu, lty = j)
legend(x = min(mu), y = 1,legend = paste("n =",nn),lty = 1:length(nn),bty = "n")

```



Sample size based on desired power

To find the smallest sample size guaranteeing a desired power:

1. Fix an alternative value μ^* and a desired power γ^* .
2. Set up the equation $\gamma(\mu^*) = \gamma^*$ and solve for n (then round up).

For our tests concerning μ when σ is known, we obtain:

► In the one-sided case $n = \left\lceil \sigma^2 \left(\frac{z_\alpha + z_{\beta^*}}{\mu^* - \mu_0} \right)^2 \right\rceil$.

► In the two-sided case $n = \left\lceil \sigma^2 \left(\frac{z_{\alpha/2} + z_{\beta^*}}{\mu^* - \mu_0} \right)^2 \right\rceil$.

In the above $\beta^* = 1 - \gamma^*$.

Exercise: Derive the sample size formula for the test of

$$H_0: \mu \leq \mu_0 \text{ vs } H_1: \mu > \mu_0$$

when σ is known.

Golden ratio example (cont):

Suppose the true mean of B/A in the population is 1.65.

Give the sample size n required to reject $H_0: \mu \leq 1.618$ vs $H_1: \mu > 1.618$ with power ≥ 0.80 . Use $S_n = 0.148$ as a guess of σ .

```
alpha <- 0.05
gm <- 0.80
sigma <- sd(gr)
mu <- 1.65
mu0 <- 1.618
za <- qnorm(1 - alpha)
zb <- qnorm(gm)
nr <- ceiling(sigma^2 * (za + zb)^2 / (mu - mu0)^2)
nr
```

```
[1] 133
```