Probability and Statistics

Nels Grevstad

Metropolitan State University of Denver

ngrevsta@msudenver.edu

May 1, 2019

Topics

1. Large Sample Test for a General Population Mean μ
2. Test for a Normal Mean μ of When σ is Unknown

Objectives

Objectives:

- Carry out a hypothesis test for a general population mean μ when the sample size is large.
- Carry out a hypothesis test for a normal mean μ when the standard deviation σ is unknown.

Large Sample Test for a General Population Mean μ

(8.2)

- When the **sample size** n **is large**, two things are true:

Large Sample Test for a General Population Mean μ

(8.2)

- When the **sample size** n **is large**, two things are true:

  1. Regardless of the shape of the population, by the Central Limit Theorem,

$$Z = \frac{\bar{X} - \mu}{\sigma/\sqrt{n}} \sim \text{N}(0, 1) \qquad \text{(approximately)}$$

Large Sample Test for a General Population Mean μ (8.2)

- When the **sample size n is large**, two things are true:

1. Regardless of the shape of the population, by the Central Limit Theorem,

$$Z = \frac{\bar{X} - \mu}{\sigma/\sqrt{n}} \sim \text{N}(0, 1) \qquad \text{(approximately)}$$

2. The **sample standard deviation S** will remain fairly **constant** from one sample to the next, and **approximately equal to σ.**

- As a consequence, **when n is large**,

$$\frac{\bar{X} - \mu}{S/\sqrt{n}} \sim \mathrm{N}(0,\, 1) \qquad \text{(approximately)}$$

even if the sample is from **non-normal** population.

(Note that σ was replaced by S above.)

- As a consequence, **when n is large**,

$$\frac{\bar{X} - \mu}{S/\sqrt{n}} \sim \text{N}(0, 1) \qquad \text{(approximately)}$$

even if the sample is from **non-normal** population.

(Note that σ was replaced by S above.)

- It follows that **if H_0 is true** (so $\mu = \mu_0$),

$$\frac{\bar{X} - \mu_0}{S/\sqrt{n}} \sim \text{N}(0, 1) \qquad \text{(approximately)}$$

- As a consequence, **when n is large**,

$$\frac{\bar{X} - \mu}{S/\sqrt{n}} \sim \text{N}(0,\, 1) \qquad \text{(approximately)}$$

  even if the sample is from **non-normal** population.

  (Note that $\boldsymbol{\sigma}$ was replaced by S above.)

- It follows that **if $\boldsymbol{H_0}$ is true** (so $\mu = \mu_0$),

$$\frac{\bar{X} - \mu_0}{S/\sqrt{n}} \sim \text{N}(0,\, 1) \qquad \text{(approximately)}$$

- Thus we can use this as our **test statistic** in a ***one-sample z test for μ***.

One-Sample Z Test Statistic:

$$Z = \frac{\bar{X} - \mu_0}{S/\sqrt{n}}.$$

When H_0 is true, $Z \sim$ N$(0, 1)$.

Sampling Distribution of the Test Statistic Under H_0:
If Z is the one-sample Z test statistic (from the previous slide), then when

$$H_0 : \ \mu \ = \ \mu_0$$

is true,

$$Z \ \sim \ \mathsf{N}(0, 1).$$

Sampling Distribution of the Test Statistic Under H_0:
If Z is the one-sample Z test statistic (from the previous slide), then when

$$H_0 : \mu = \mu_0$$

is true,

$$Z \sim \text{N}(0, 1).$$

- The ***p-value*** is the appropriate tail area under the N$(0,\ 1)$ curve. (See Slides 17.)

Sampling Distribution of the Test Statistic Under H_0:
If Z is the one-sample Z test statistic (from the previous slide), then when

$$H_0: \ \mu \ = \ \mu_0$$

is true,

$$Z \sim \text{N}(0, 1).$$

- The ***p-value*** is the appropriate tail area under the N$(0,\ 1)$ curve. (See Slides 17.)
- In practice, n is **large enough** if $n > 40$.

One-Sample Z Test for μ when n is Large

Assumptions: The data $x_1, x_2, \ldots, x_n$ are a random sample from *any* distribution whose mean is μ and n is large.

Null hypothesis: $H_0 : \mu = \mu_0$.

Test statistic value: $z = \frac{\bar{x} - \mu_0}{s/\sqrt{n}}$.

Decision rule: Reject H_0 if p-value $< \alpha$.

Alternative hypothesis	**P-value** = area under $\mathrm{N}(0, 1)$ distribution:
$H_a : \mu > \mu_0$	to the right of z
$H_a : \mu < \mu_0$	to the left of z
$H_a : \mu \neq \mu_0$	to the left of $-\lvert z \rvert$ and right of $\lvert z \rvert$

Test for a Normal Mean μ of When σ is Unknown (8.3)

- Recall (Slides 15) that if $X_1, X_2, \ldots, X_n$ are a random sample from a **normal** population, the random variable

$$\frac{\bar{X} - \mu}{S/\sqrt{n}} \sim t(n-1),$$

a t ***distribution*** with $n-1$ ***degrees of freedom***.

Test for a Normal Mean μ of When σ is Unknown (8.3)

- Recall (Slides 15) that if $X_1, X_2, \ldots, X_n$ are a random sample from a **normal** population, the random variable

$$\frac{\bar{X} - \mu}{S/\sqrt{n}} \sim t(n-1),$$

a t ***distribution*** with $n - 1$ ***degrees of freedom***.

- It follows that **if H_0 is true** (so $\mu = \mu_0$),

$$\frac{\bar{X} - \mu_0}{S/\sqrt{n}} \sim t(n-1).$$

- Thus we can use this as our **test statistic** in a ***one-sample*** t ***test for*** μ.

- The **test statistic** for the ***one-sample*** t ***test for*** μ is

One-Sample t Test Statistic:

$$T = \frac{\bar{X} - \mu_0}{S/\sqrt{n}}$$

When H_0 is true, $T \sim t(n-1)$.

Sampling Distribution of the Test Statistic Under H_0:
If T is the one-sample t test statistic, then when

$$H_0: \ \mu \ = \ \mu_0$$

is true,

$$T \ \sim \ \mathrm{t}(n-1).$$

Sampling Distribution of the Test Statistic Under H_0:
If T is the one-sample t test statistic, then when

$$H_0: \ \mu \ = \ \mu_0$$

is true,

$$T \ \sim \ \text{t}(n-1).$$

- The ***p-value*** is the appropriate tail area under the $\text{t}(n-1)$ curve. (See the next few slides.)

P-Value: For the **one-sample t test**, the **p-value** is the tail area under the $t(n-1)$ curve:

1. To the **right** of the observed t if the alternative hypothesis is $H_a : \mu > \mu_0$:

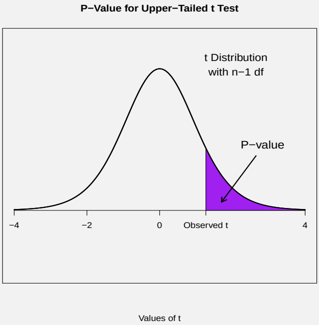

1. To the **left** of the observed t if the alternative hypothesis is $H_a : \mu < \mu_0$:

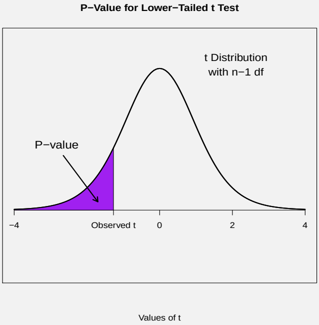

1. To the **left** of $-|t|$ **and right** of $|t|$ if the alternative hypothesis is $H_a : \mu \neq \mu_0$:

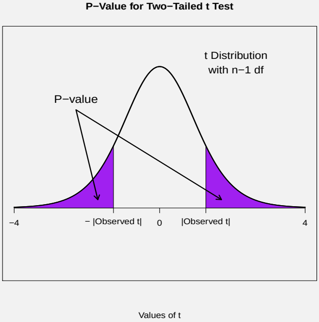

One-Sample t Test for μ

Assumptions: The data $x_1, x_2, \ldots, x_n$ are a random sample from a *normal* distribution.

Null hypothesis: $H_0 : \mu = \mu_0$.

Test statistic value: $t = \frac{\bar{x} - \mu_0}{s/\sqrt{n}}$.

Decision rule: Reject H_0 if p-value $< \alpha$.

Alternative hypothesis	P-value = area under the t distribution with $n-1$ d.f.:
$H_a : \mu > \mu_0$	to the right of t
$H_a : \mu < \mu_0$	to the left of t
$H_a : \mu \neq \mu_0$	to the left of $-\lvert t \rvert$ and right of $\lvert t \rvert$

Example

A quality control engineer monitors a machine that puts cereal into boxes.

Example

A quality control engineer monitors a machine that puts cereal into boxes.

According to the label, each box is supposed to contain **16** oz of cereal.

Example

A quality control engineer monitors a machine that puts cereal into boxes.

According to the label, each box is supposed to contain **16** oz of cereal.

The machine will need to be adjusted if the boxes are systematically being **under-filled** or **over-filled**.

Example

A quality control engineer monitors a machine that puts cereal into boxes.

According to the label, each box is supposed to contain **16** oz of cereal.

The machine will need to be adjusted if the boxes are systematically being **under-filled** or **over-filled**.

From past experience, the engineer knows that the weight (ounces) of the cereal in a box follows a **normal** distribution.

To decide if the boxes are being **under-filled or overfilled**, the engineer will test the **hypotheses**

$$\begin{aligned} H_0 : \mu &= 16 \\ H_a : \mu &\neq 16 \end{aligned}$$

where μ is the true (unknown) population mean weight.

To decide if the boxes are being **under-filled or overfilled**, the engineer will test the **hypotheses**

$$H_0: \mu = 16$$
$$H_a: \mu \neq 16$$

where μ is the true (unknown) population mean weight.

A random sample of **ten** boxes gives

$$\bar{x} = \mathbf{16.6} \qquad \text{and} \qquad s = \mathbf{0.9}.$$

Thus the observed **test statistic** is

$$t = \frac{\bar{x} - \mu_0}{s/\sqrt{n}}$$

Thus the observed **test statistic** is

$$t = \frac{\bar{x} - \mu_0}{s/\sqrt{n}} = \frac{16.6-16}{0.9/\sqrt{10}}$$

Thus the observed **test statistic** is

$$\begin{aligned} t &= \frac{\bar{x} - \mu_0}{s/\sqrt{n}} \\ &= \frac{16.6 - 16}{0.9/\sqrt{10}} \\ &= \mathbf{2.11}. \end{aligned}$$

Thus the **sample mean** weight, $\bar{x} = 16.6$, is about **2.11 standard errors above 16** ounces.

Thus the observed **test statistic** is

$$t = \frac{\bar{x} - \mu_0}{s/\sqrt{n}} = \frac{16.6 - 16}{0.9/\sqrt{10}} = \mathbf{2.11}.$$

Thus the **sample mean** weight, $\bar{x} = \mathbf{16.6}$, is about **2.11 standard errors above 16** ounces.

The **p-value** is the **probability** that we'd get a t value this far away from zero (in either direction) by chance **if** the **population mean** μ was **16**.

From the **two tail** areas of the **sampling distribution** that the test statistic would follow under H_0 (the t(9) distribution), to the **right** of **2.11** and **left** of **-2.11**,

$$\textbf{p-value} \;=\; 2(0.033) \;=\; \mathbf{0.066}.$$

Thus we'd get a result like the one we got **6.6%** of the time **even if** the **population mean** μ was **16** ounces.

Thus we'd get a result like the one we got **6.6%** of the time **even if** the **population mean** μ was **16** ounces.

Using a **level of significance** $\alpha = 0.05$, the **decision rule** is

Reject H_0 if p-value < 0.05.
Fail to reject H_0 if p-value ≥ 0.05.

Thus we'd get a result like the one we got **6.6%** of the time **even if** the **population mean** μ was **16** ounces.

Using a **level of significance** $\alpha = 0.05$, the **decision rule** is

Reject H_0 if p-value < 0.05.
Fail to reject H_0 if p-value ≥ 0.05.

Because $0.066 \geq 0.05$, we **fail to reject** H_0.

Thus we'd get a result like the one we got **6.6%** of the time **even if** the **population mean** μ was **16** ounces.

Using a **level of significance** $\alpha = 0.05$, the **decision rule** is

Reject H_0 if p-value < 0.05.
Fail to reject H_0 if p-value ≥ 0.05.

Because $\mathbf{0.066 \geq 0.05}$, we **fail to reject** $\boldsymbol{H_0}$.

There's **no statistically significant evidence** that the population mean cereal box weight μ is different from 16 ounces.

Thus we'd get a result like the one we got **6.6%** of the time **even if** the **population mean** μ was **16** ounces.

Using a **level of significance** $\alpha = 0.05$, the **decision rule** is

Reject H_0 if p-value < 0.05.
Fail to reject H_0 if p-value ≥ 0.05.

Because $0.066 \geq 0.05$, we **fail to reject H_0.**

There's **no statistically significant evidence** that the population mean cereal box weight μ is different from 16 ounces.

The result that the engineer got (by taking a random sample) can be explained by chance variation (sampling error).