Induction Course in Quantitative Methods for Finance

Estimation: Single Population

Definitions

- An estimator of a population parameter is
  - a random variable that depends on sample information . . .
  - whose value provides an approximation to this unknown parameter
- A specific value of that random variable is called an estimate

Point and Interval Estimates

- A point estimate is a single number,
- a confidence interval provides additional information about variability

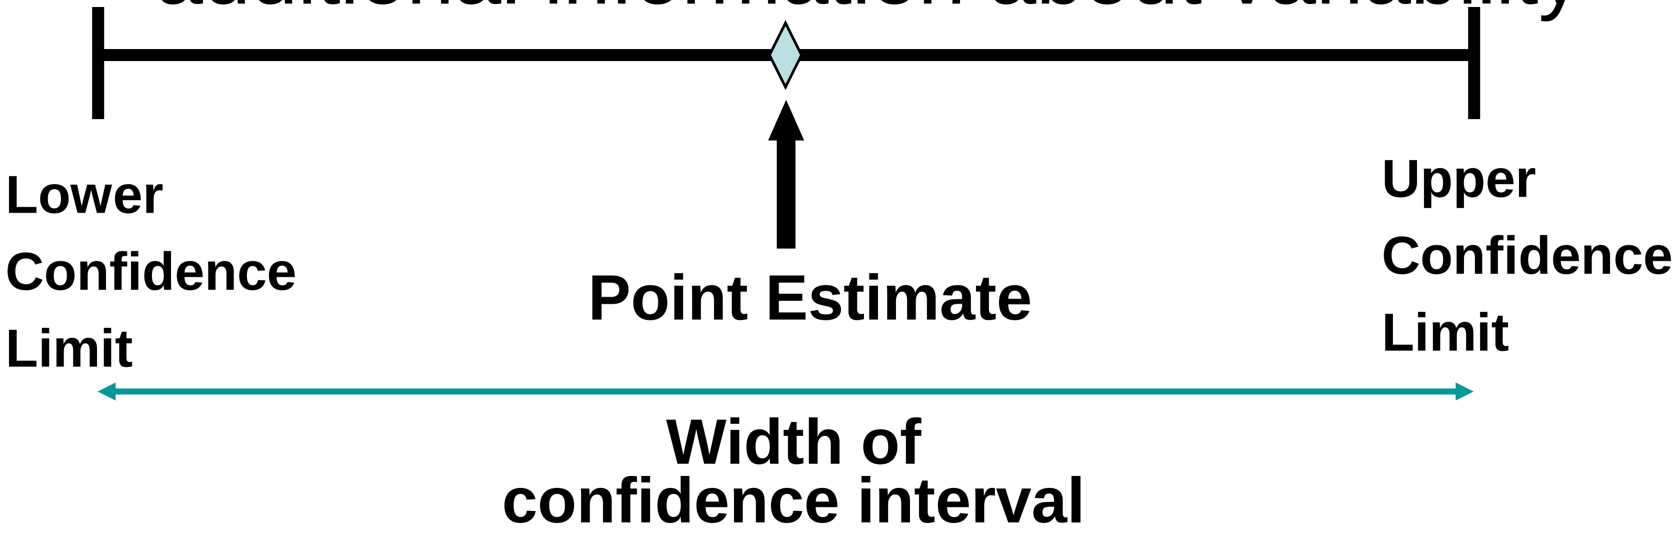

Point Estimates

We can estimate a Population Parameter …		with a Sample Statistic (a Point Estimate)
Mean	μ	$\bar{x}$
Proportion	P	$\hat{p}$

Unbiasedness

- A point estimator $\hat{\theta}$ is said to be an unbiased estimator of the parameter θ if the expected value, or mean, of the sampling distribution of $\hat{\theta}$ is θ,

$$E(\hat{\theta}) = \theta$$

- Examples:
  - The sample mean is an unbiased estimator of μ
  - The sample variance is an unbiased estimator of σ^2
  - The sample proportion is an unbiased estimator of P

Unbiasedness

(continued)

- $\hat{\theta}_1$ is an unbiased estimator, $\hat{\theta}_2$ is biased:

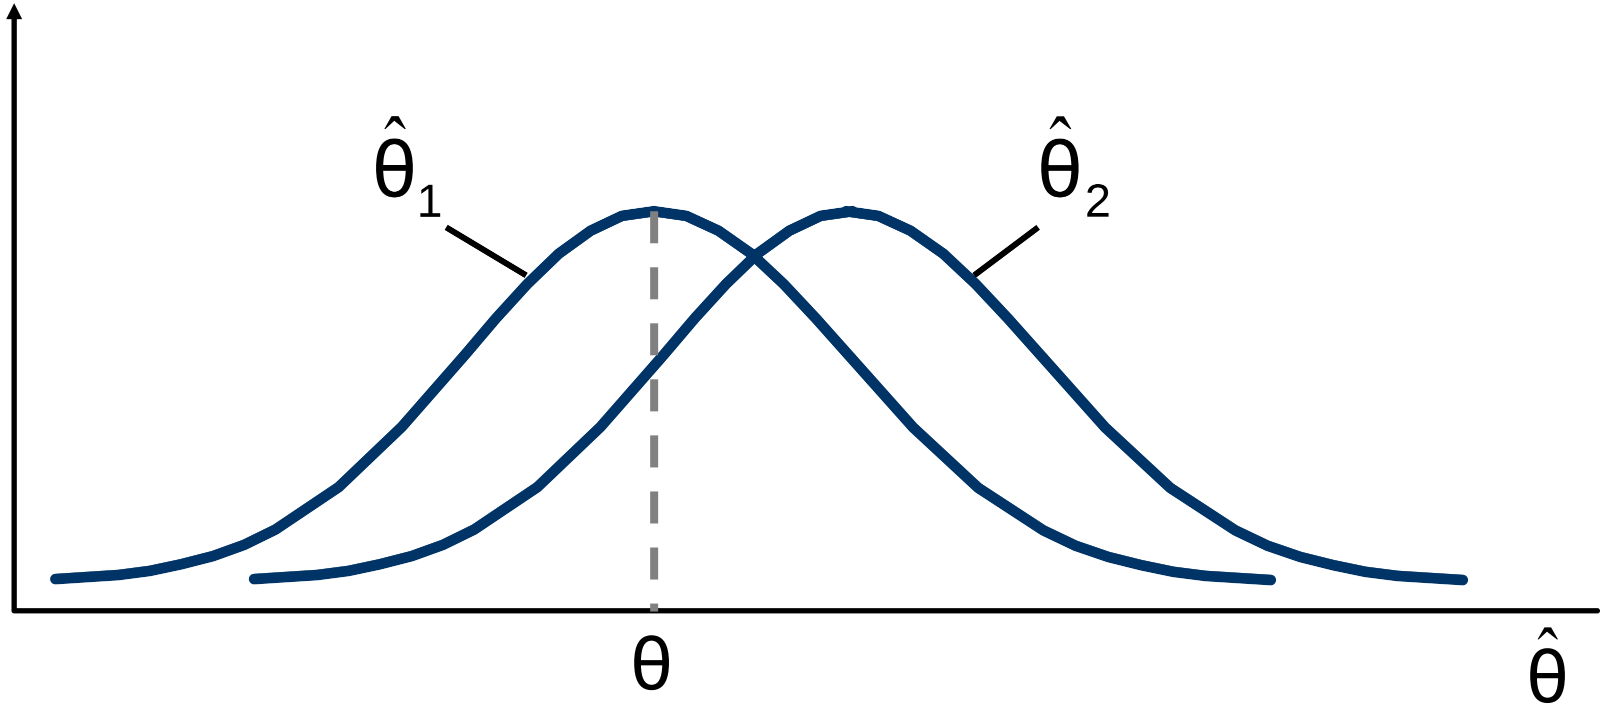

Bias

- Let $\hat{\theta}$ be an estimator of θ
- The bias in $\hat{\theta}$ is defined as the difference between its mean and θ

$$\text{Bias}(\hat{\theta}) = E(\hat{\theta}) - \theta$$

- The bias of an unbiased estimator is 0

Consistency

- Let $\hat{\theta}$ be an estimator of θ
- $\hat{\theta}$ is a consistent estimator of θ if the difference between the expected value of and θ decreases as the sample size increases

- Consistency is desired when unbiased estimators cannot be obtained

Most Efficient Estimator

- Suppose there are several unbiased estimators of θ
- The most efficient estimator or the minimum variance unbiased estimator of θ is the unbiased estimator with the smallest variance

- Let $\hat{\theta}_1$ and $\hat{\theta}_2$ be two unbiased estimators of θ, based on the same number of sample observations. Then,
  - $\hat{\theta}_1$ is said to be more efficient than $\hat{\theta}_2$ if $\text{Var}(\hat{\theta}_1) < \text{Var}(\hat{\theta}_2)$
  - The relative efficiency of $\hat{\theta}_1$ with respect to $\hat{\theta}_2$ is the ratio of their variances:

$$\text{Relative Efficiency} = \frac{\text{Var}(\hat{\theta}_2)}{\text{Var}(\hat{\theta}_1)}$$

Confidence Intervals

- How much uncertainty is associated with a point estimate of a population parameter?
- An interval estimate provides more information about a population characteristic than does a point estimate
- Such interval estimates are called confidence intervals

Confidence Interval Estimate

- An interval gives a range of values:
  - Takes into consideration variation in sample statistics from sample to sample
  - Based on observation from 1 sample
  - Gives information about closeness to unknown population parameters
  - Stated in terms of level of confidence
    - Can never be 100% confident

Confidence Interval and Confidence Level

- If $P(a < \theta < b) = 1 - \alpha$ then the interval from a to b is called a $100(1 - \alpha)\%$ confidence interval of θ.

- The quantity $(1-\alpha)$ is called the confidence level of the interval (α between 0 and 1)

  – In repeated samples of the population, the true value of the parameter θ would be contained in $100(1 - \alpha)\%$ of intervals calculated this way.
  – The confidence interval calculated in this manner is written as $a < \theta < b$ with $100(1 - \alpha)\%$ confidence

Confidence Level, (1-α)

(continued)

- Suppose confidence level = 95%
- Also written $(1 - \alpha) = 0.95$
- A relative frequency interpretation:
  - From repeated samples, 95% of all the confidence intervals that can be constructed will contain the unknown true parameter
- A specific interval either will contain or will not contain the true parameter
  - No probability involved in a specific interval

General Formula

- The general formula for all confidence intervals is:

Point Estimate ± (Reliability Factor)(Standard Error)

- The value of the reliability factor depends on the desired level of confidence

Confidence Interval for μ (σ^2 Known)

- Assumptions
  - Population variance σ^2 is known
  - Population is normally distributed
  - If population is not normal, use large sample
- Confidence interval estimate:

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

(where $z_{\alpha/2}$ is the normal distribution value for a probability of $\alpha/2$ in each tail)

Margin of Error

- The confidence interval,

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

- Can also be written as $\bar{x} \pm ME$

  where ME is called the margin of error

$$ME = z_{\alpha/2}\frac{\sigma}{\sqrt{n}}$$

- The interval width, w, is equal to twice the margin of error

Reducing the Margin of Error

$$ME = z_{\alpha/2} \frac{\sigma}{\sqrt{n}}$$

The margin of error can be reduced if

- the population standard deviation can be reduced ($\sigma\downarrow$)
- The sample size is increased ($n\uparrow$)
- The confidence level is decreased, $(1-\alpha)\downarrow$

Finding the Reliability Factor, $z_{\alpha/2}$

- Consider a 95% confidence interval:

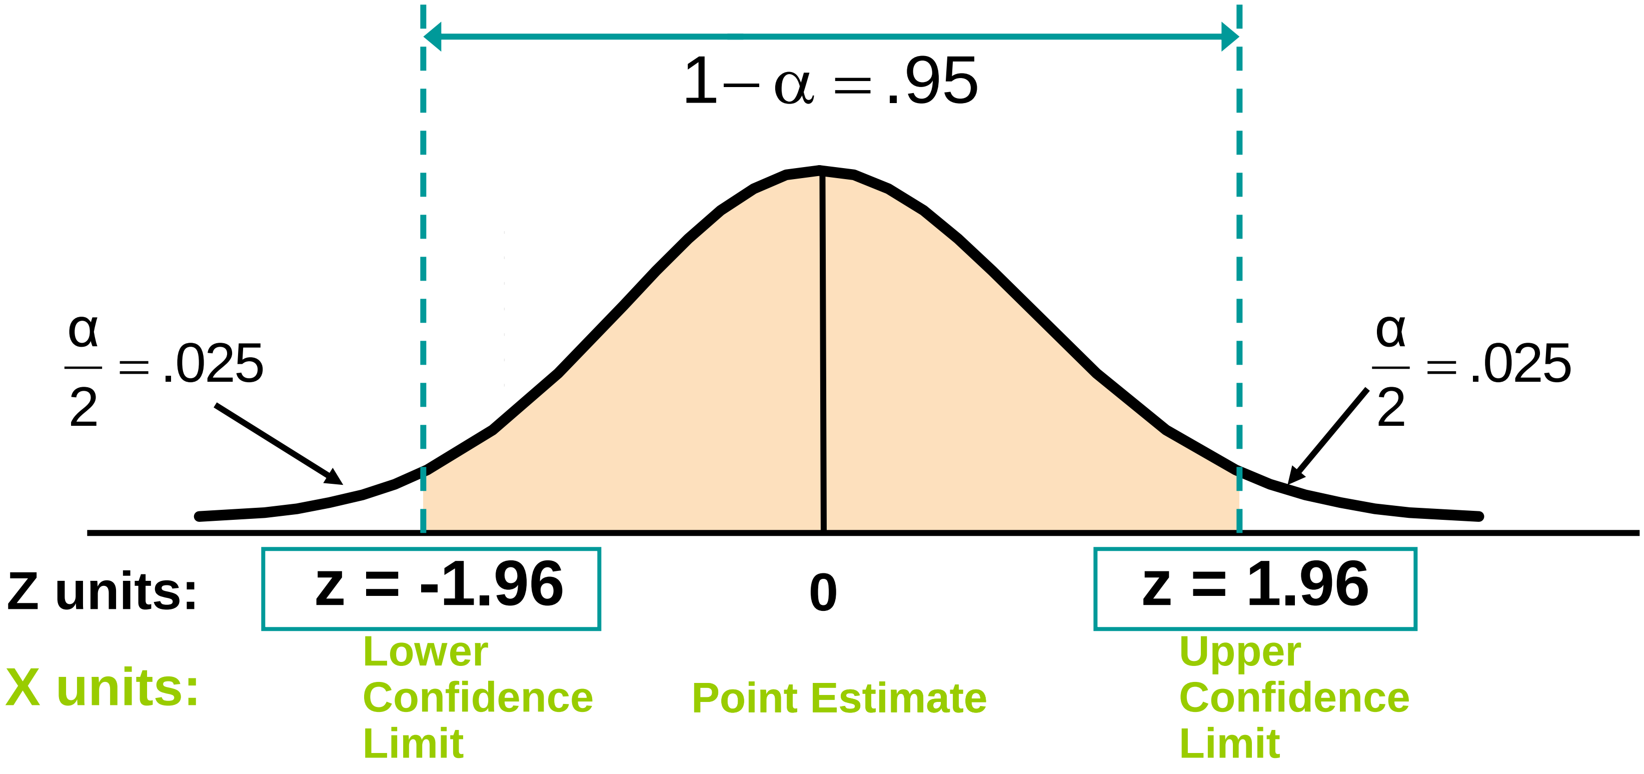

- Find $z_{.025} = \pm 1.96$ from the standard normal distribution table

Common Levels of Confidence

- Commonly used confidence levels are 90%, 95%, and 99%

Confidence Level	***Confidence Coefficient,*** $1-\alpha$	$Z_{\alpha/2}$ ***value***
80%	**.80**	**1.28**
90%	**.90**	**1.645**
95%	**.95**	**1.96**
98%	**.98**	**2.33**
99%	**.99**	**2.58**
99.8%	**.998**	**3.08**
99.9%	**.999**	**3.27**

Intervals and Level of Confidence

Sampling Distribution of the Mean

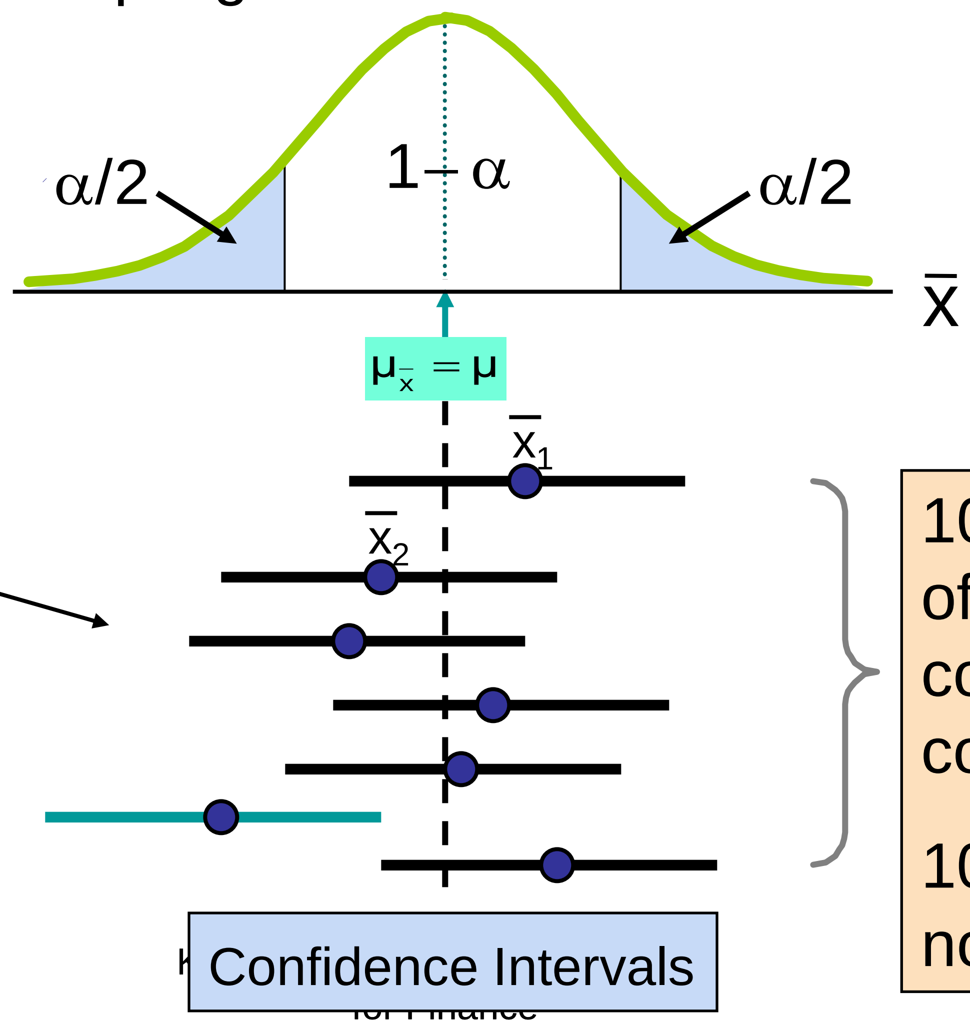

Intervals extend from

$$\bar{x} - z\frac{\sigma}{\sqrt{n}}$$

to

$$\bar{x} + z\frac{\sigma}{\sqrt{n}}$$

100(1-α)% of intervals constructed contain μ;

100(α)% do not.

Example

- A sample of 11 firms from a large normal population has a mean monthly return of 2.20%. We know from past testing that the population standard deviation is 0.35%.

- Determine a 95% confidence interval for the true mean return of the population.

Example

(continued)

$$\bar{x} \pm z \frac{\sigma}{\sqrt{n}}$$

$$= 2.20 \pm 1.96\,(.35/\sqrt{11})$$

$$= 2.20 \pm .2068$$

$$1.9932 \;<\; \mu \;<\; 2.4068$$

Interpretation

- We are 95% confident that the true mean return is between 1.9932 and 2.4068 %
- Although the true mean may or may not be in this interval, 95% of intervals formed in this manner will contain the true mean

Student's t Distribution

- Consider a random sample of n observations
  - with mean $\bar{x}$ and standard deviation s
  - from a normally distributed population with mean μ

- Then the variable $$t = \frac{\bar{x} - \mu}{s/\sqrt{n}}$$

- follows the Student's t distribution with (n - 1) degrees of freedom

Confidence Interval for μ (σ^2 Unknown)

- If the population standard deviation σ is unknown, we can substitute the sample standard deviation, s
- This introduces extra uncertainty, since s is variable from sample to sample
- So we use the t distribution instead of the normal distribution

Confidence Interval for μ (σ Unknown)

(continued)

- Assumptions
  - Population standard deviation is unknown
  - Population is normally distributed
  - If population is not normal, use large sample
- Use Student's t Distribution

- Confidence Interval Estimate:

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

where $t_{n-1,\alpha/2}$ is the critical value of the t distribution with n-1 d.f. and an area of $\alpha/2$ in each tail:

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

Student's t Distribution

- The t is a family of distributions
- The t value depends on degrees of freedom (d.f.)
  - Number of observations that are free to vary after sample mean has been calculated

$$d.f. = n - 1$$

Student's t Distribution

Note: t ⟶ Z as n increases

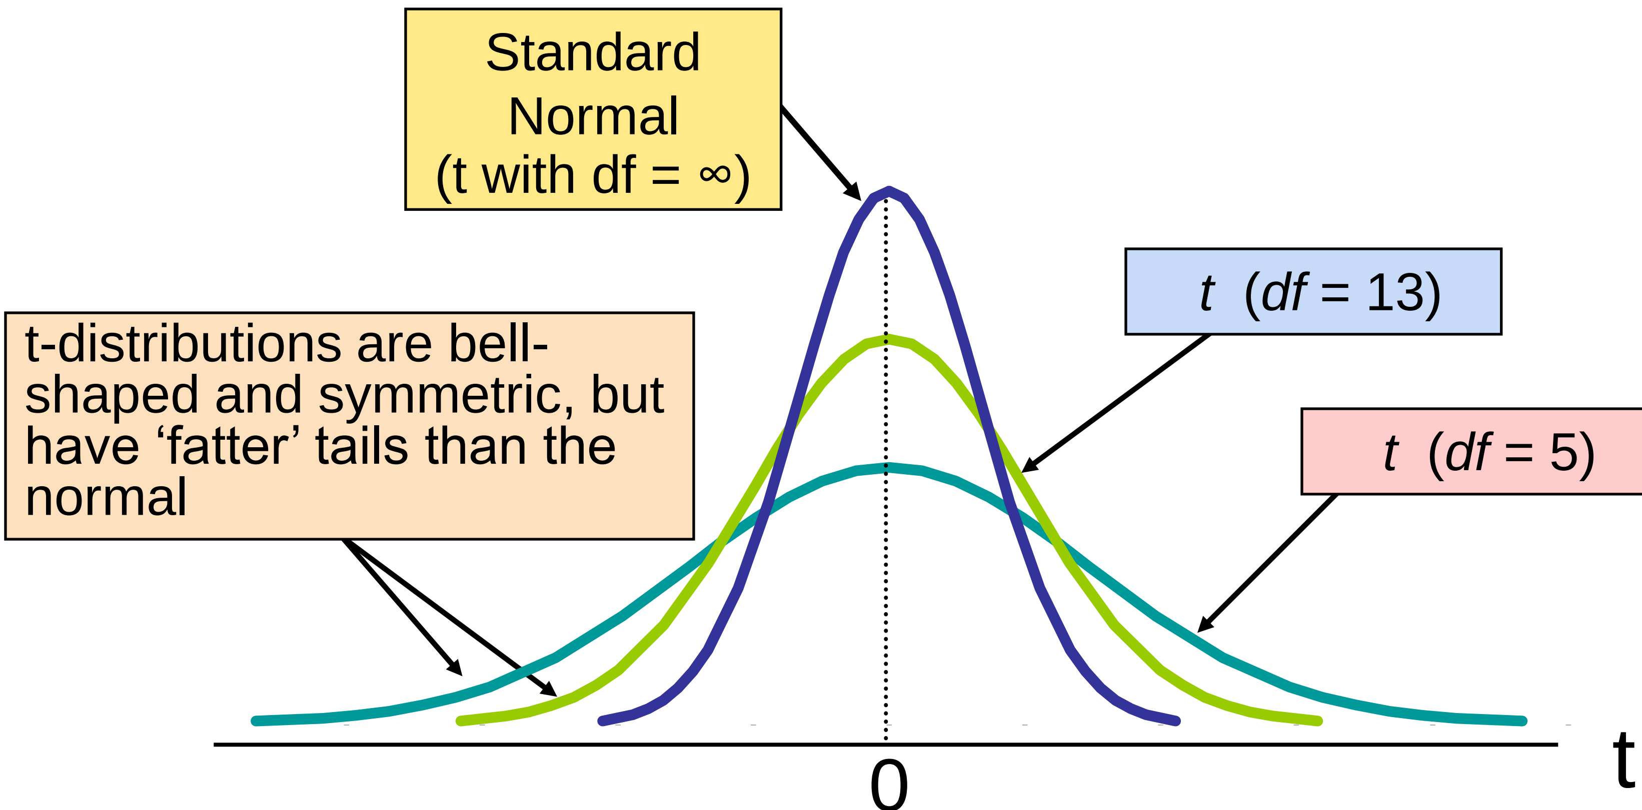

Student's t Table

df	Upper Tail Area .10	.05	.025
1	3.078	6.314	12.706
2	1.886	2.920	4.303
3	1.638	2.353	3.182

The body of the table contains t values, not probabilities

Let: n = 3

df = n - 1 = 2

α = .10

α/2 =.05

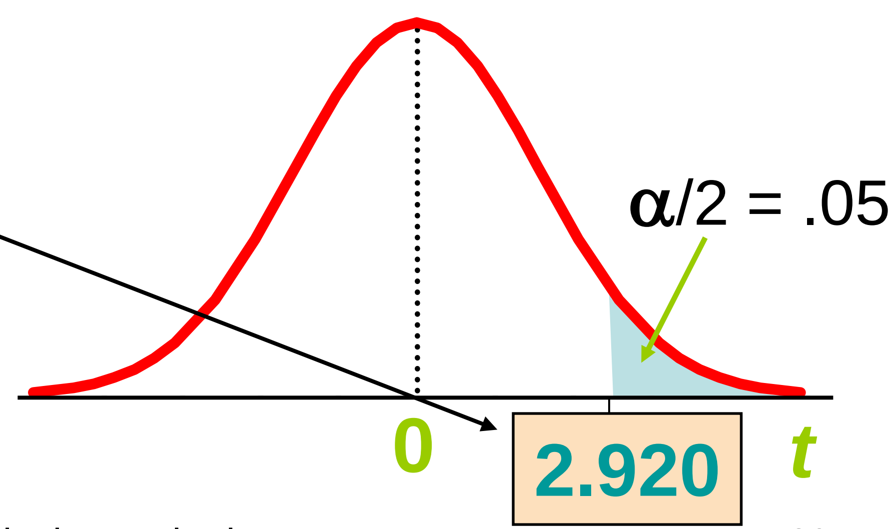

t distribution values

With comparison to the Z value

Confidence Level	t (10 d.f.)	t (20 d.f.)	t (30 d.f.)	Z
.80	1.372	1.325	1.310	1.282
.90	1.812	1.725	1.697	1.645
.95	2.228	2.086	2.042	1.960
.99	3.169	2.845	2.750	2.576

Note: $t \longrightarrow Z$ as n increases

Example

A random sample of n = 25 has $\bar{x}$ = 50 and s = 8. Form a 95% confidence interval for μ

– d.f. = n – 1 = 24, so $t_{n-1,\alpha/2} = t_{24,.025} = 2.0639$

The confidence interval is

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

$$50 - (2.0639)\frac{8}{\sqrt{25}} < \mu < 50 + (2.0639)\frac{8}{\sqrt{25}}$$

$$46.698 < \mu < 53.302$$

Confidence Intervals for the Population Proportion, p

- An interval estimate for the population proportion (P) can be calculated by adding an allowance for uncertainty to the sample proportion ($\hat{p}$)

Confidence Intervals for the Population Proportion, p

(continued)

- Recall that the distribution of the sample proportion is approximately normal if the sample size is large, with standard deviation

$$\sigma_P = \sqrt{\frac{P(1-P)}{n}}$$

- We will estimate this with sample data:

$$\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$$

Confidence Interval Endpoints

- Upper and lower confidence limits for the population proportion are calculated with the formula

$$\hat{p} - z_{\alpha/2}\sqrt{\frac{\hat{p}(1-\hat{p})}{n}} < P < \hat{p} + z_{\alpha/2}\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$$

- where
  - $z_{\alpha/2}$ is the standard normal value for the level of confidence desired
  - $\hat{p}$ is the sample proportion
  - n is the sample size

Example

- A random sample of 100 firms shows that 25 did not pay dividend
- Form a 95% confidence interval for the true proportion of non-paying firms

Example

(continued)

$$\hat{p} - z_{\alpha/2}\sqrt{\frac{\hat{p}(1-\hat{p})}{n}} \; < \; P \; < \; \hat{p} + z_{\alpha/2}\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$$

$$\frac{25}{100} - 1.96\sqrt{\frac{.25(.75)}{100}} \; < \; P \; < \; \frac{25}{100} + 1.96\sqrt{\frac{.25(.75)}{100}}$$

$$0.1651 \; < \; P \; < \; 0.3349$$

Interpretation

- We are 95% confident that the true percentage of non-paying firms in the population is between

  16.51% and 33.49%.

- Although the interval from 0.1651 to 0.3349 may or may not contain the true proportion, 95% of intervals formed from samples of size 100 in this manner will contain the true proportion.

Dependent Samples

Dependent samples

Tests Means of 2 Related Populations

- Paired or matched samples
- Repeated measures (before/after)
- Use difference between paired values:

$$d_i = x_i - y_i$$

- Assumptions:
  - Both Populations Are Normally Distributed

Mean Difference

Dependent samples

The i^{th} paired difference is d_i , where

$$d_i = x_i - y_i$$

The point estimate for the population mean paired difference is $\bar{d}$:

$$\bar{d} = \frac{\sum_{i=1}^{n} d_i}{n}$$

The sample standard deviation is:

$$S_d = \sqrt{\frac{\sum_{i=1}^{n} (d_i - \bar{d})^2}{n-1}}$$

n is the number of matched pairs in the sample

Confidence Interval for Mean Difference

Dependent samples

The confidence interval for difference between population means, μ_d , is

$$\bar{d} - t_{n-1,\alpha/2} \frac{S_d}{\sqrt{n}} < \mu_d < \bar{d} + t_{n-1,\alpha/2} \frac{S_d}{\sqrt{n}}$$

Where

n = the sample size

(number of matched pairs in the paired sample)

Confidence Interval for Mean Difference

(continued)

Dependent samples

- The margin of error is

$$ME = t_{n-1,\alpha/2} \frac{s_d}{\sqrt{n}}$$

- $t_{n-1,\alpha/2}$ is the value from the Student's t distribution with (n – 1) degrees of freedom for which

$$P(t_{n-1} > t_{n-1,\alpha/2}) = \frac{\alpha}{2}$$

Paired Samples Example

- Six people sign up for a weight loss program. You collect the following data:

Person	Weight: Before (x)	After (y)	Difference, d_i
1	136	125	11
2	205	195	10
3	157	150	7
4	138	140	- 2
5	175	165	10
6	166	160	6
			42

$$\bar{d} = \frac{\sum d_i}{n}$$

$$= 7.0$$

$$S_d = \sqrt{\frac{\sum (d_i - \bar{d})^2}{n-1}}$$

$$= 4.82$$

Paired Samples Example

(continued)

- For a 95% confidence level, the appropriate t value is $t_{n-1,\alpha/2} = t_{5,.025} = 2.571$
- The 95% confidence interval for the difference between means, μ_d , is

$$\bar{d} - t_{n-1,\alpha/2} \frac{S_d}{\sqrt{n}} < \mu_d < \bar{d} + t_{n-1,\alpha/2} \frac{S_d}{\sqrt{n}}$$

$$7 - (2.571) \frac{4.82}{\sqrt{6}} < \mu_d < 7 + (2.571) \frac{4.82}{\sqrt{6}}$$

$$-1.94 < \mu_d < 12.06$$

Since this interval contains zero, we cannot be 95% confident, given this limited data, that the weight loss program helps people lose weight

Difference Between Two Means

Population means, independent samples

Goal: Form a confidence interval for the difference between two population means, $\mu_x - \mu_y$

- Different data sources
  - Unrelated
  - Independent
    - Sample selected from one population has no effect on the sample selected from the other population
- The point estimate is the difference between the two sample means:

$$\bar{x} - \bar{y}$$

Difference Between Two Means

(continued)

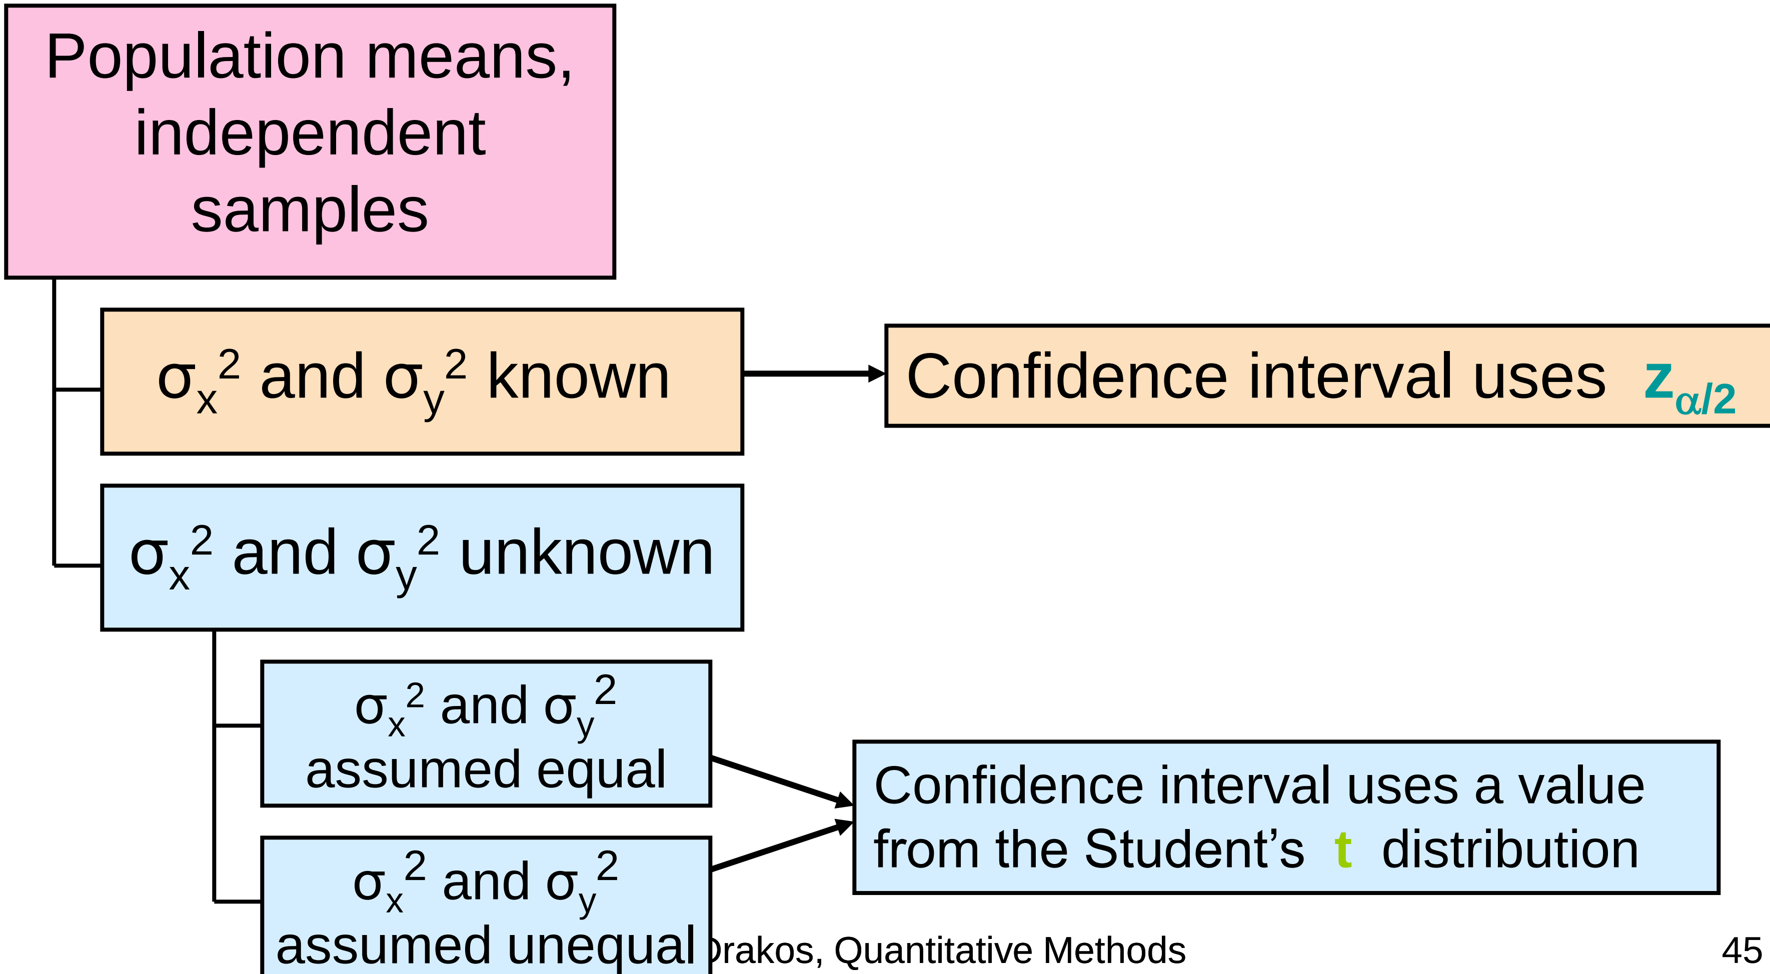

σ_x^2 and σ_y^2 Known

Population means, independent samples

- σ_x^2 and σ_y^2 known *
- σ_x^2 and σ_y^2 unknown

Assumptions:

- Samples are randomly and independently drawn
- both population distributions are normal
- Population variances are known

σ_x^2 and σ_y^2 Known

(continued)

Population means, independent samples

- σ_x^2 and σ_y^2 known *
- σ_x^2 and σ_y^2 unknown

When σ_x and σ_y are known and both populations are normal, the variance of $\overline{X} - \overline{Y}$ is

$$\sigma^2_{\overline{X}-\overline{Y}} = \frac{\sigma_x^2}{n_x} + \frac{\sigma_y^2}{n_y}$$

…and the random variable

$$Z = \frac{(\overline{x} - \overline{y}) - (\mu_X - \mu_Y)}{\sqrt{\dfrac{\sigma_x^2}{n_X} + \dfrac{\sigma_y^2}{n_Y}}}$$

has a standard normal distribution

Confidence Interval, σ_x^2 and σ_y^2 Known

Population means, independent samples

- σ_x^2 and σ_y^2 known
- σ_x^2 and σ_y^2 unknown

* The confidence interval for $\mu_x - \mu_y$ is:

$$(\bar{x}-\bar{y}) - z_{\alpha/2}\sqrt{\frac{\sigma_X^2}{n_x}+\frac{\sigma_Y^2}{n_y}} < \mu_X - \mu_Y < (\bar{x}-\bar{y}) + z_{\alpha/2}\sqrt{\frac{\sigma_X^2}{n_x}+\frac{\sigma_Y^2}{n_y}}$$

σ_x^2 and σ_y^2 Unknown, Assumed Equal

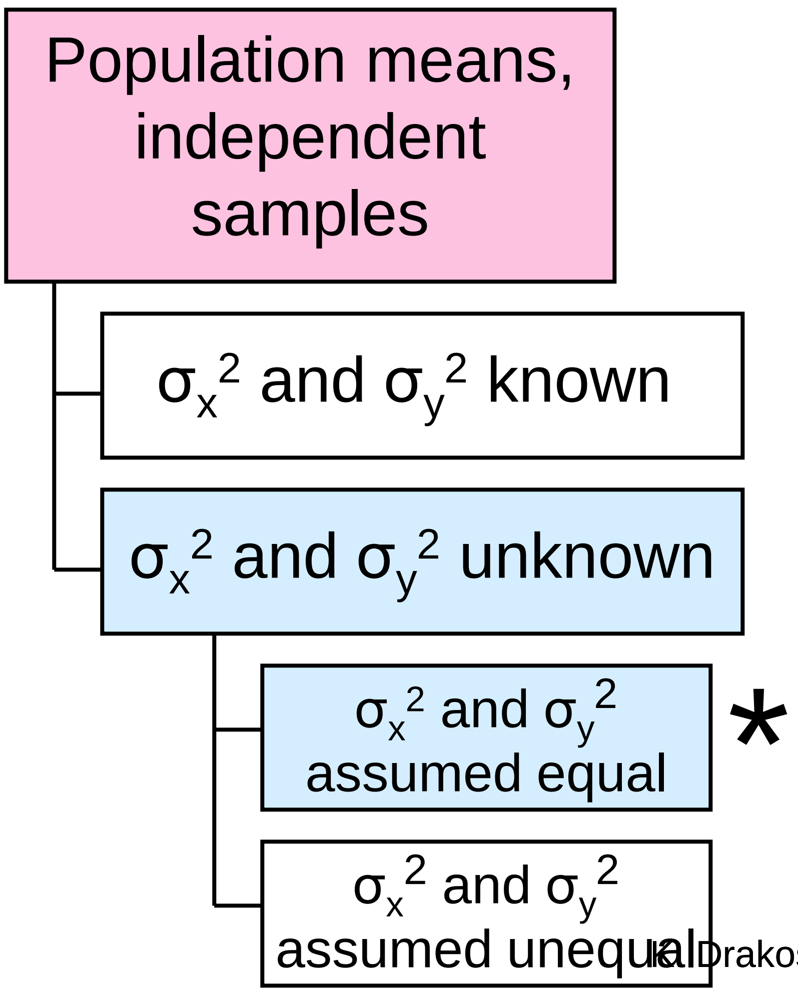

Assumptions:

- Samples are randomly and independently drawn
- Populations are normally distributed
- Population variances are unknown but assumed equal

σ_x2 and σ_y^2 Unknown, Assumed Equal

(continued)

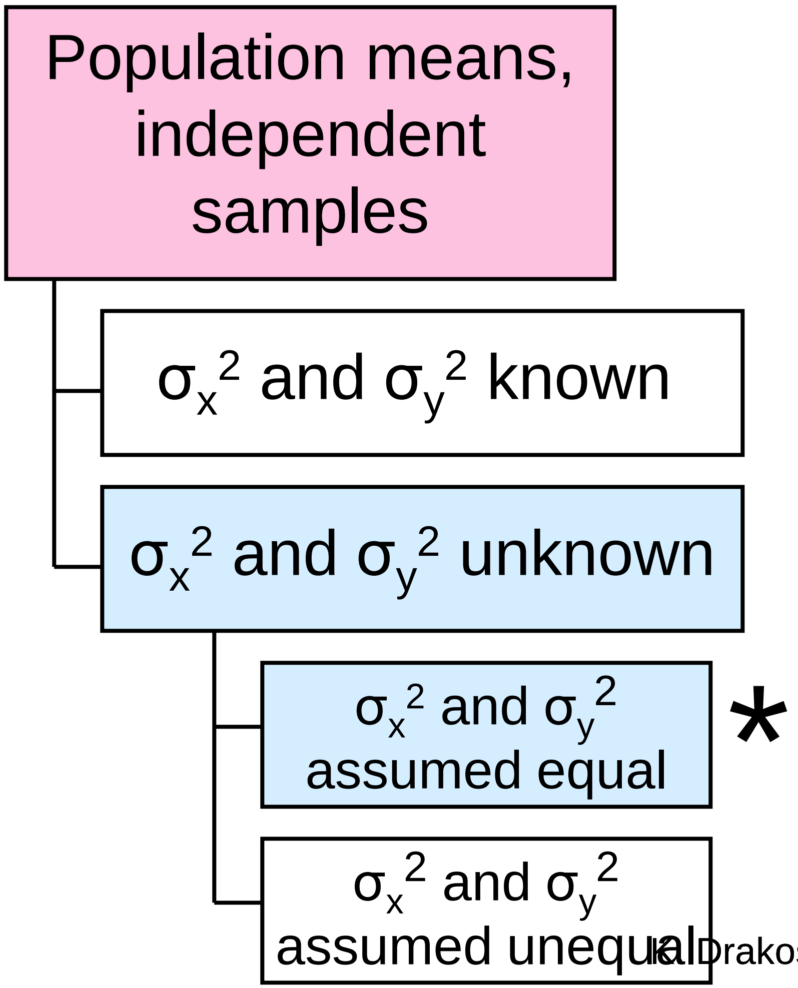

Forming interval estimates:

- The population variances are assumed equal, so use the two sample standard deviations and pool them to estimate σ
- use a t value with $(n_x + n_y – 2)$ degrees of freedom

σ_x^2 and σ_y^2 Unknown, Assumed Equal

(continued)

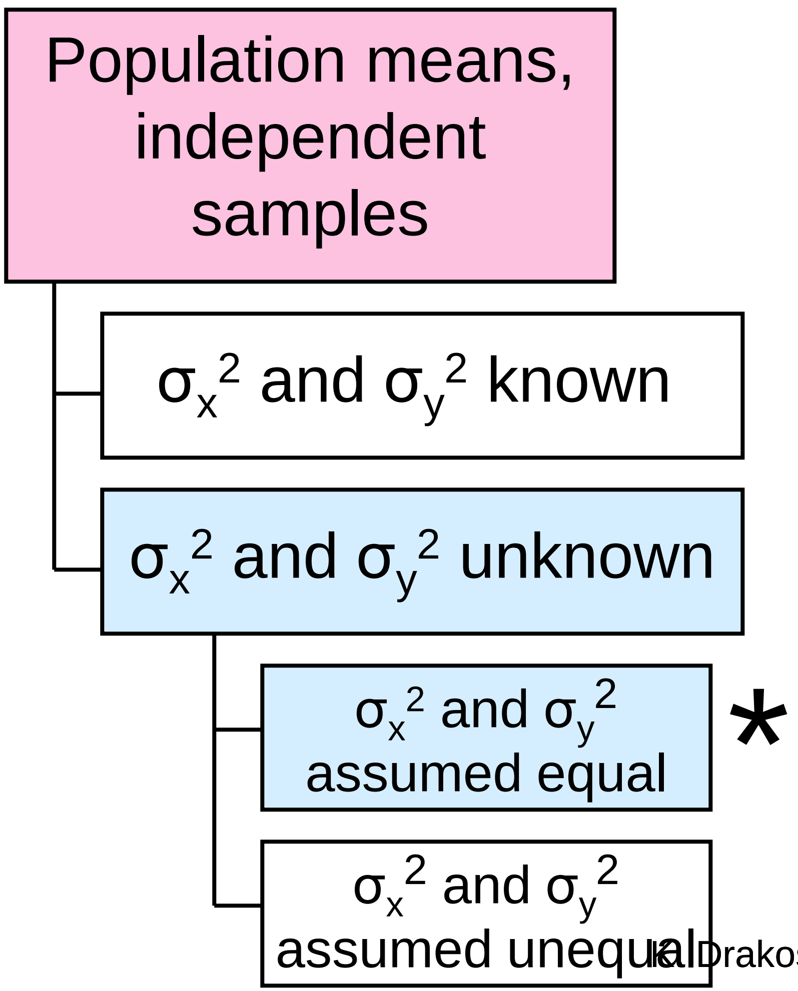

The pooled variance is

$$s_p^2 = \frac{(n_x - 1)s_x^2 + (n_y - 1)s_y^2}{n_x + n_y - 2}$$

Confidence Interval, σ_x^2 and σ_y^2 Unknown, Equal

- σ_x^2 and σ_y^2 unknown
  - σ_x^2 and σ_y^2 assumed equal *
  - σ_x^2 and σ_y^2 assumed unequal

The confidence interval for $\mu_1 - \mu_2$ is:

$$(\bar{x}-\bar{y}) - t_{n_x+n_y-2,\alpha/2}\sqrt{\frac{s_p^2}{n_x}+\frac{s_p^2}{n_y}} < \mu_X - \mu_Y < (\bar{x}-\bar{y}) + t_{n_x+n_y-2,\alpha/2}\sqrt{\frac{s_p^2}{n_x}+\frac{s_p^2}{n_y}}$$

Where

$$s_p^2 = \frac{(n_x-1)s_x^2 + (n_y-1)s_y^2}{n_x+n_y-2}$$

σ_x^2 and σ_y^2 Unknown, Assumed Unequal

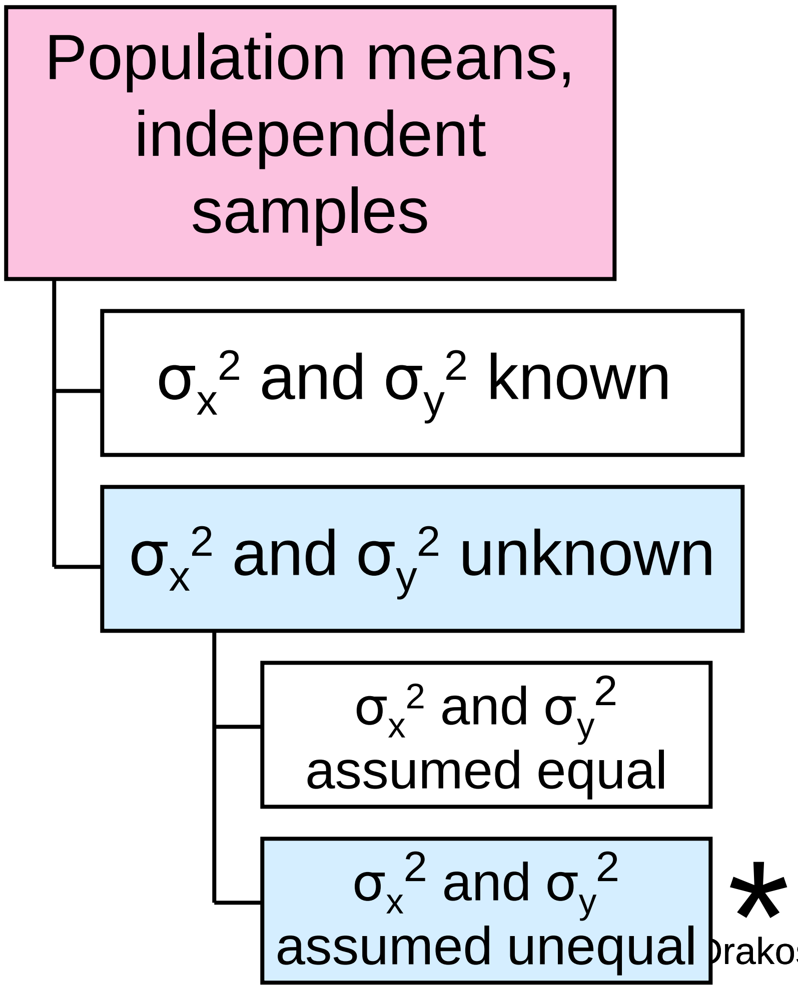

Assumptions:

- Samples are randomly and independently drawn
- Populations are normally distributed
- Population variances are unknown and assumed unequal

σ_x^2 and σ_y^2 Unknown, Assumed Unequal

(continued)

- Population means, independent samples
  - σ_x^2 and σ_y^2 known
  - σ_x^2 and σ_y^2 unknown
    - σ_x^2 and σ_y^2 assumed equal
    - σ_x^2 and σ_y^2 assumed unequal *

Forming interval estimates:

- The population variances are assumed unequal, so a pooled variance is not appropriate
- use a t value with ν degrees of freedom, where

$$\nu = \frac{\left[\left(\frac{s_x^2}{n_x}\right) + \left(\frac{s_y^2}{n_y}\right)\right]^2}{\left(\frac{s_x^2}{n_x}\right)^2 / (n_x - 1) + \left(\frac{s_y^2}{n_y}\right)^2 / (n_y - 1)}$$

Confidence Interval, σ_x^2 and σ_y^2 Unknown, Unequal

σ_x^2 and σ_y^2 unknown

- σ_x^2 and σ_y^2 assumed equal
- σ_x^2 and σ_y^2 assumed unequal *

The confidence interval for $\mu_1 - \mu_2$ is:

$$(\bar{x}-\bar{y}) - t_{\nu,\alpha/2}\sqrt{\frac{s_x^2}{n_x}+\frac{s_y^2}{n_y}} < \mu_X - \mu_Y < (\bar{x}-\bar{y}) + t_{\nu,\alpha/2}\sqrt{\frac{s_x^2}{n_x}+\frac{s_y^2}{n_y}}$$

Where

$$\nu = \frac{\left[\left(\frac{s_x^2}{n_x}\right)+\left(\frac{s_y^2}{n_y}\right)\right]^2}{\left(\frac{s_x^2}{n_x}\right)^2/(n_x-1)+\left(\frac{s_y^2}{n_y}\right)^2/(n_y-1)}$$

Two Population Proportions

Population proportions

Goal: Form a confidence interval for the difference between two population proportions, $P_x - P_y$

Assumptions:

Both sample sizes are large (generally at least 40 observations in each sample)

The point estimate for the difference is $\hat{p}_x - \hat{p}_y$

Two Population Proportions

(continued)

Population proportions

- The random variable

$$Z = \frac{(\hat{p}_x - \hat{p}_y) - (p_x - p_y)}{\sqrt{\frac{\hat{p}_x(1-\hat{p}_x)}{n_x} + \frac{\hat{p}_y(1-\hat{p}_y)}{n_y}}}$$

is approximately normally distributed

Confidence Interval for Two Population Proportions

Population proportions

The confidence limits for $P_x - P_y$ are:

$$(\hat{p}_x - \hat{p}_y) \pm Z_{\alpha/2}\sqrt{\frac{\hat{p}_x(1-\hat{p}_x)}{n_x} + \frac{\hat{p}_y(1-\hat{p}_y)}{n_y}}$$

Example:
Two Population Proportions

Form a 90% confidence interval for the difference between the proportion of retail firms and the proportion of industrial firms who went bankrupt last year.

- In a random sample, 26 of 50 retail and 28 of 40 industrial firms had gone bankrupt

Example: Two Population Proportions

(continued)

Retail:

$$\hat{p}_x = \frac{26}{50} = 0.52$$

Industrial:

$$\hat{p}_y = \frac{28}{40} = 0.70$$

$$\sqrt{\frac{\hat{p}_x(1-\hat{p}_x)}{n_x} + \frac{\hat{p}_y(1-\hat{p}_y)}{n_y}} = \sqrt{\frac{0.52(0.48)}{50} + \frac{0.70(0.30)}{40}} = 0.1012$$

For 90% confidence, $Z_{\alpha/2} = 1.645$

Example: Two Population Proportions

(continued)

The confidence limits are:

$$(\hat{p}_x - \hat{p}_y) \pm Z_{\alpha/2} \sqrt{\frac{\hat{p}_x(1-\hat{p}_x)}{n_x} + \frac{\hat{p}_y(1-\hat{p}_y)}{n_y}}$$

$$= (.52 - .70) \pm 1.645(0.1012)$$

so the confidence interval is

$$-0.3465 < P_x - P_y < -0.0135$$

Since this interval does not contain zero we are 90% confident that the two proportions are not equal

Confidence Intervals for the Population Variance

Population Variance

- Goal: Form a confidence interval for the population variance, σ^2

- The confidence interval is based on the sample variance, s^2

- Assumed: the population is normally distributed

Confidence Intervals for the Population Variance

(continued)

Population Variance

The random variable

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

follows a chi-square distribution with (n – 1) degrees of freedom

The chi-square value $\chi^2_{n-1,\alpha}$ denotes the number for which

$$P(\chi^2_{n-1} > \chi^2_{n-1,\alpha}) = \alpha$$

Confidence Intervals for the Population Variance

(continued)

Population Variance

The (1 - α)% confidence interval for the population variance is

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

Hypothesis Testing

What is a Hypothesis?

- A hypothesis is a claim (assumption) about a population parameter:
- population mean / population proportion

Example: The mean monthly cell phone bill of this city is $\mu = \$42$

Example: The proportion of adults in this city with cell phones is $p = .68$

The Null Hypothesis, H_0

- States the assumption (numerical) to be tested

> Example: The average number of TV sets in U.S. Homes is equal to three $(H_0 : \mu = 3)$

- Is always about a population parameter, not about a sample statistic

$$H_0 : \mu = 3$$

$$H_0 : \bar{X} = 3$$

The Null Hypothesis, H_0

(continued)

- Begin with the assumption that the null hypothesis is true
  - Similar to the notion of innocent until proven guilty
- Refers to the status quo
- Always contains "=" , "≤" or "≥" sign
- May or may not be rejected

The Alternative Hypothesis, H_1

- Is the opposite of the null hypothesis
  - e.g., The average number of TV sets in U.S. homes is not equal to 3 ($H_1: \mu \neq 3$)
- Challenges the status quo
- May or may not be supported
- Is generally the hypothesis that the researcher is trying to support

Hypothesis Testing Process

Claim: the population mean age is 50. (Null Hypothesis: $H_0: \mu = 50$)

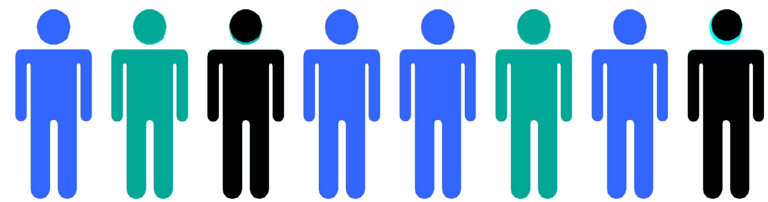

Population

Now select a random sample

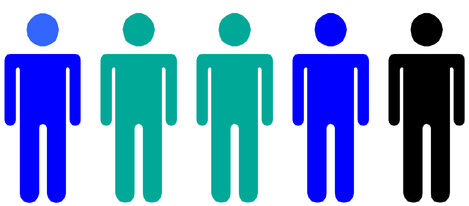

Sample

Suppose the sample mean age is 20: $\bar{X} = 20$

Is $\bar{X} = 20$ likely if $\mu = 50$?

If not likely,

REJECT

Null Hypothesis

Hypothesis Tests Design

- Is $\overline{X}$ likely given that $\mu = 50$? If we believe that this not likely we will reject H_0.
- How can we determine if the event is likely to occur given that H_0 is true?
- We define a *rejection region* of the sampling distribution, $\overline{X} < c$.

$$P\left(\overline{X} < c \middle| \mu = 50\right) = P\left(\text{Reject } H_0 \middle| H_0 \text{ true}\right) = \alpha$$

- So we want small values of α (*significance level*). According to α if we know the distribution of X we can determine c (*critical value*)

Hypothesis Tests Design

(continued)

- If we find that for $\alpha = 1\%$ the $c = 22$ then since $20 < 22$ *we will reject H_0 at the 1% significance level*.
- So for H_0 being true the 1% of the samples would have $\overline{X} < c$. The rest 99% would have a sample mean $\overline{X} > c$. So we are 99% confident that H_0 should be rejected.

Hypothesis Tests Design

(continued)

Sampling Distribution of $\bar{X}$

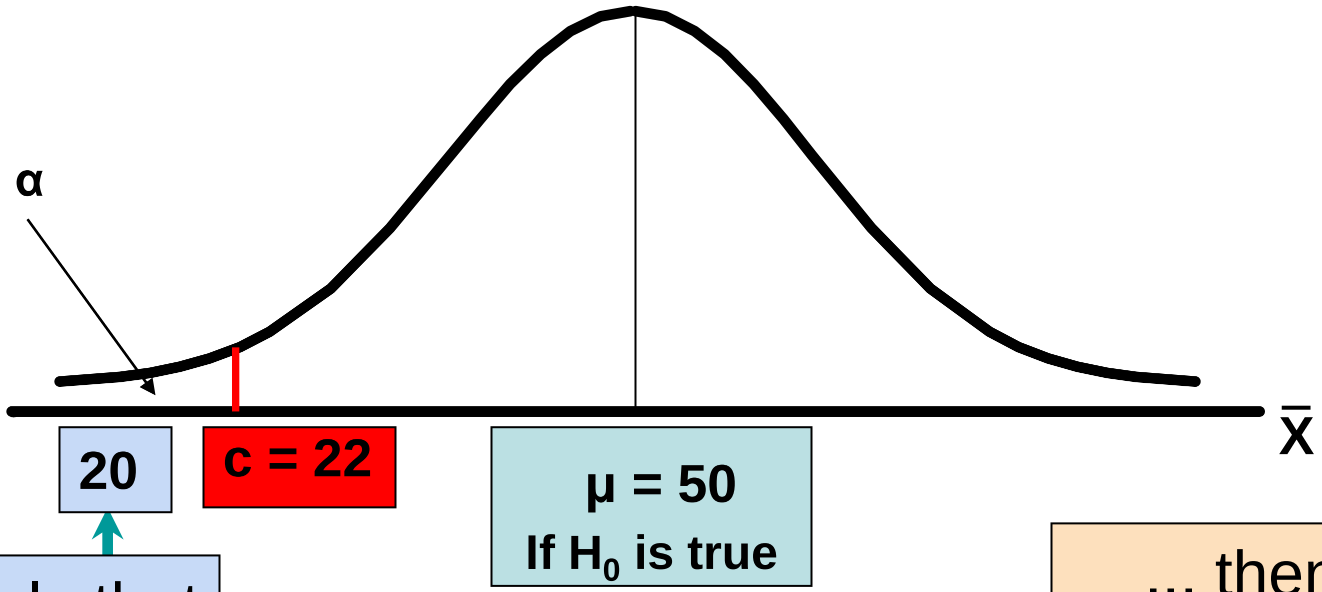

If it is unlikely that we would get a sample mean of this value ...

... if in fact this were the population mean…

... then we reject the null hypothesis that μ = 50.

Level of Significance, α

- Defines the unlikely values of the sample statistic if the null hypothesis is true
  - Defines rejection region of the sampling distribution
- Is designated by α , (level of significance)
  - Typical values are .01, .05, or .10
- Is selected by the researcher at the beginning
- Provides the critical value(s) of the test

Level of Significance and the Rejection Region

Level of significance = α

Represents critical value

Rejection region is shaded

H_0: $\mu = 3$

H_1: $\mu \neq 3$

Two-tail test

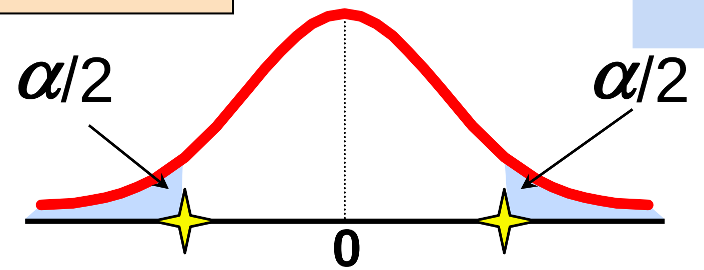

H_0: $\mu \leq 3$

H_1: $\mu > 3$

Upper-tail test

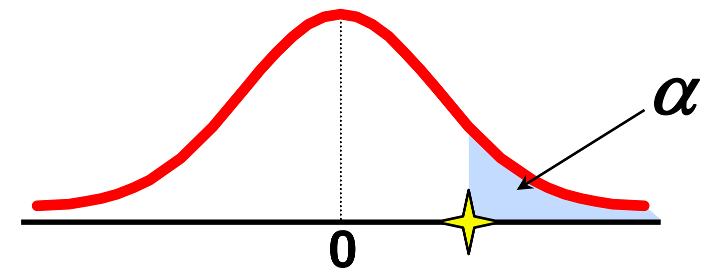

H_0: $\mu \geq 3$

H_1: $\mu < 3$

Lower-tail test

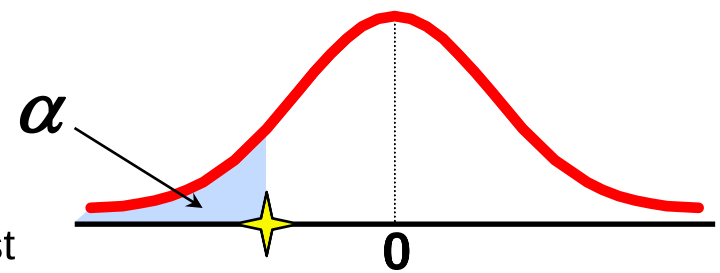

Errors in Making Decisions

- **Type I Error**
  - Reject a true null hypothesis
  - Considered a serious type of error

The probability of Type I Error is α

- Called level of significance of the test
- Set by researcher in advance

Errors in Making Decisions

(continued)

- **Type II Error**
  - Fail to reject a false null hypothesis

The probability of Type II Error is β

Outcomes and Probabilities

Possible Hypothesis Test Outcomes

	Actual Situation	
Decision	H_0 True	H_0 False
Do Not Reject $\mathbf{H_0}$	**No error** **(1 - α)**	**Type II Error** **(β)**
Reject $\mathbf{H_0}$	**Type I Error** **(α)**	**No Error** **(1 - β)**

Key:
Outcome
(Probability)

Type I & II Error Relationship

- Type I and Type II errors can not happen at the same time
  - Type I error can only occur if H_0 is true
  - Type II error can only occur if H_0 is false

If Type I error probability (α) ↑, then

Type II error probability (β) ↓

Factors Affecting Type II Error

- All else equal,
  - β ⇧ when the difference between hypothesized parameter and its true value ⇩
  - β ⇧ when α ⇩
  - β ⇧ when σ ⇧
  - β ⇧ when n ⇩

Power of the Test

- The power of a test is the probability of rejecting a null hypothesis that is false

- i.e., Power = P(Reject H_0 | H_1 is true)

  – Power of the test increases as the sample size increases

Test of Hypothesis for the Mean (σ Known)

- Convert sample result ($\bar{x}$) to a z value

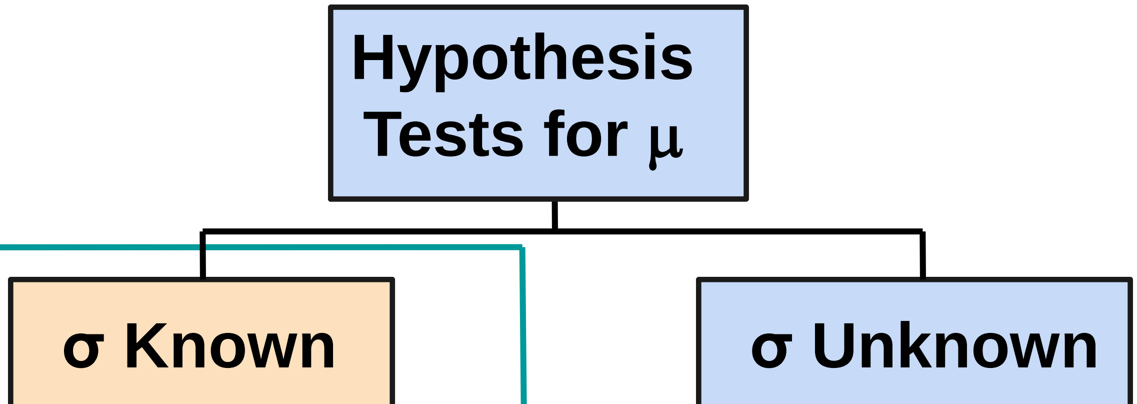

Consider the test

$$H_0 : \mu = \mu_0$$
$$H_1 : \mu > \mu_0$$

(Assume the population is normal)

The decision rule is:

$$\text{Reject } H_0 \text{ if } \quad z = \frac{\bar{x} - \mu_0}{\frac{\sigma}{\sqrt{n}}} > z_\alpha$$

Decision Rule

Reject H_0 if $z = \frac{\bar{X} - \mu_0}{\frac{\sigma}{\sqrt{n}}} > z_\alpha$

H_0: $\mu = \mu_0$

H_1: $\mu > \mu_0$

Alternate rule:

Reject H_0 if $\bar{X} > \mu_0 + Z_\alpha \sigma / \sqrt{n}$

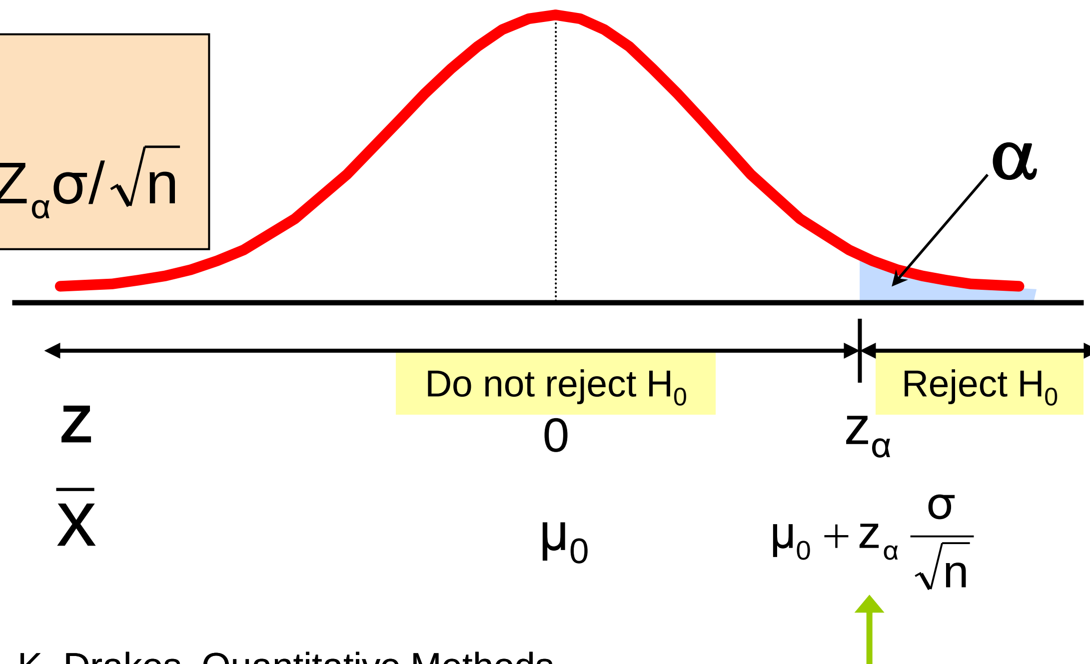

Critical value

p-Value Approach to Testing

- p-value: Probability of obtaining a test statistic more extreme ($\leq$ or $\geq$) than the observed sample value given H_0 is true
  - Also called observed level of significance
  - Smallest value of α for which H_0 can be rejected

p-Value Approach to Testing

(continued)

- Convert sample result (e.g., $\bar{x}$) to test statistic (e.g., z statistic)
- Obtain the p-value
  - For an upper tail test:

$$\text{p-value} = P\left(Z > \frac{\bar{x} - \mu_0}{\sigma/\sqrt{n}}, \text{ given that } H_0 \text{ is true}\right)$$

$$= P\left(Z > \frac{\bar{x} - \mu_0}{\sigma/\sqrt{n}} \mid \mu = \mu_0\right)$$

- Decision rule: compare the p-value to α

  - If p-value $< \alpha$, reject H_0
  - If p-value $\geq \alpha$, do not reject H_0

Example: Upper-Tail Z Test for Mean (σ Known)

A phone industry manager thinks that customer monthly cell phone bill have increased, and now average over $52 per month. The company wishes to test this claim. (Assume $\sigma = 10$ is known)

Form hypothesis test:

H_0: $\mu \leq 52$	the average is not over $52 per month
H_1: $\mu > 52$	the average **is** greater than $52 per month (i.e., sufficient evidence exists to support the manager's claim)

Example: Find Rejection Region

(continued)

- Suppose that α = .10 is chosen for this test

Find the rejection region:

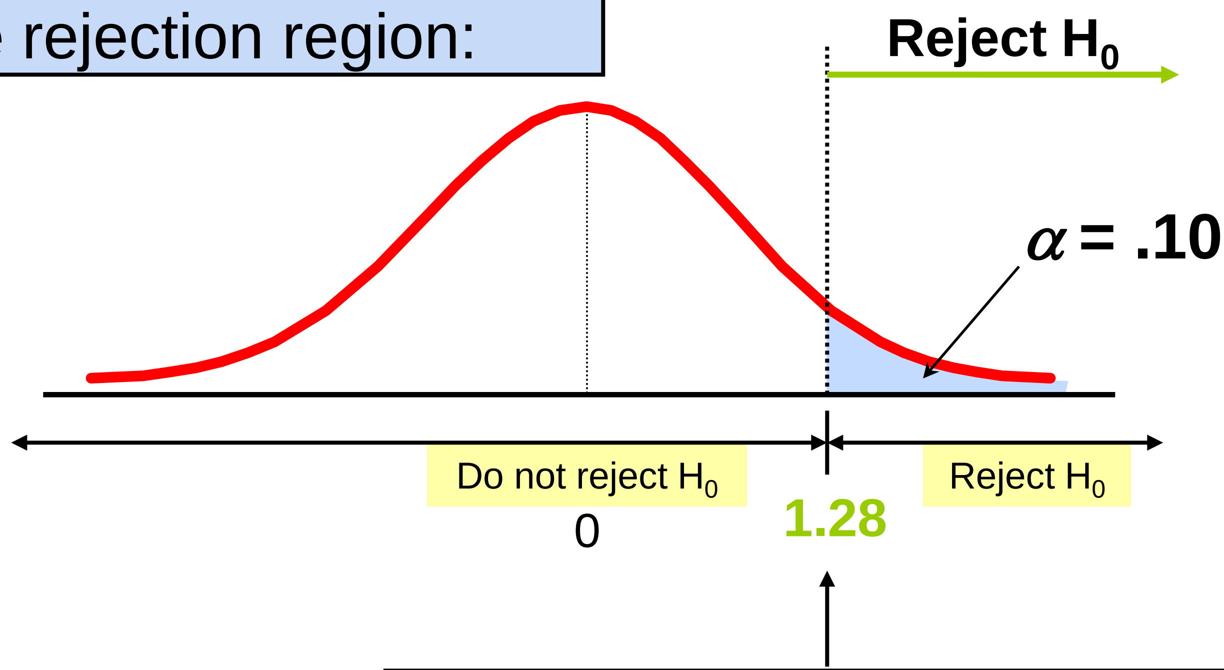

Reject H_0 if $z = \dfrac{\bar{x} - \mu_0}{\sigma / \sqrt{n}} > 1.28$

Example: Sample Results

(continued)

Obtain sample and compute the test statistic

Suppose a sample is taken with the following results: n = 64, $\bar{x}$ = 53.1 (σ=10 was assumed known)

Using the sample results,

$$z = \frac{\bar{x} - \mu_0}{\frac{\sigma}{\sqrt{n}}} = \frac{53.1 - 52}{\frac{10}{\sqrt{64}}} = 0.88$$

Example: Decision

(continued)

Reach a decision and interpret the result:

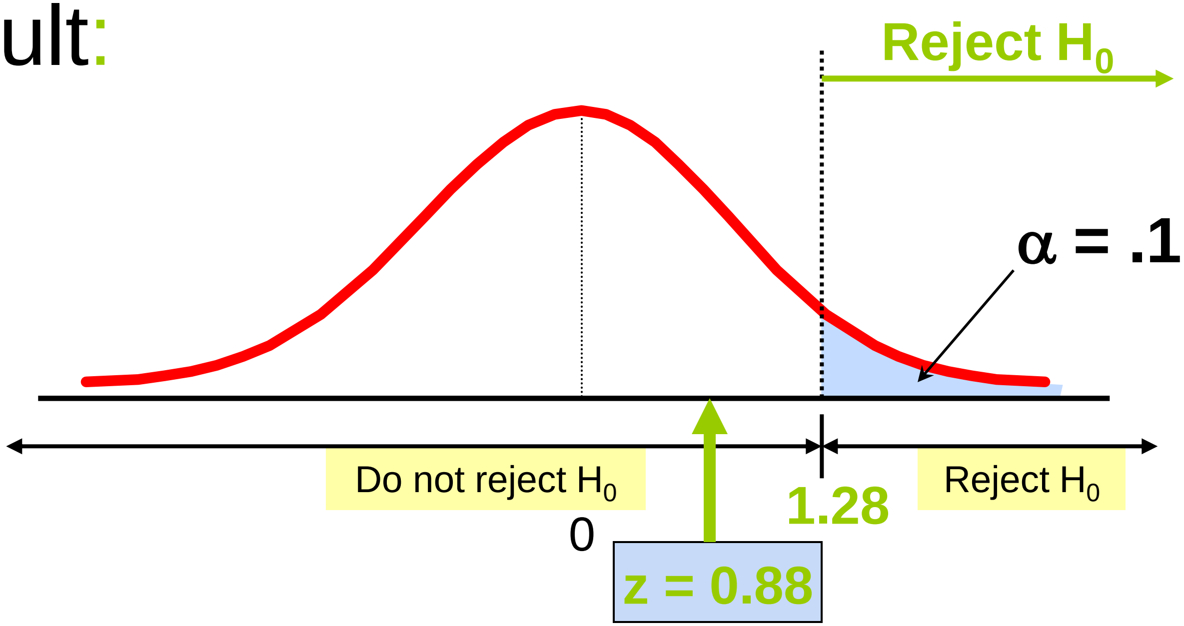

Do not reject H_0 since z = 0.88 < 1.28

i.e.: there is not sufficient evidence that the mean bill is over $52

Example: p-Value Solution

(continued)

Calculate the p-value and compare to α

(assuming that $\mu = 52.0$)

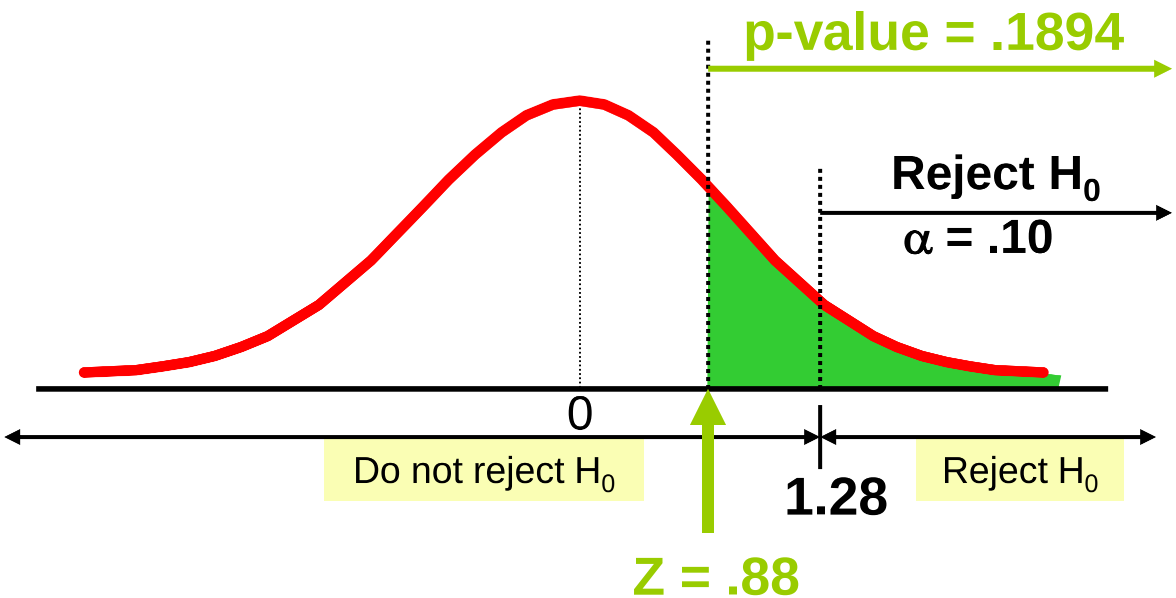

$$P(\bar{x} \geq 53.1 \mid \mu = 52.0)$$

$$= P\left(z \geq \frac{53.1 - 52.0}{10/\sqrt{64}}\right)$$

$$= P(z \geq 0.88) = 1 - .8106$$

$$= .1894$$

Do not reject H_0 since p-value = .1894 > α = .10

One-Tail Tests

- In many cases, the alternative hypothesis focuses on one particular direction

H_0: $\mu \leq 3$

H_1: $\mu > 3$

→ This is an upper-tail test since the alternative hypothesis is focused on the upper tail above the mean of 3

H_0: $\mu \geq 3$

H_1: $\mu < 3$

→ This is a lower-tail test since the alternative hypothesis is focused on the lower tail below the mean of 3

Upper-Tail Tests

- There is only one critical value, since the rejection area is in only one tail

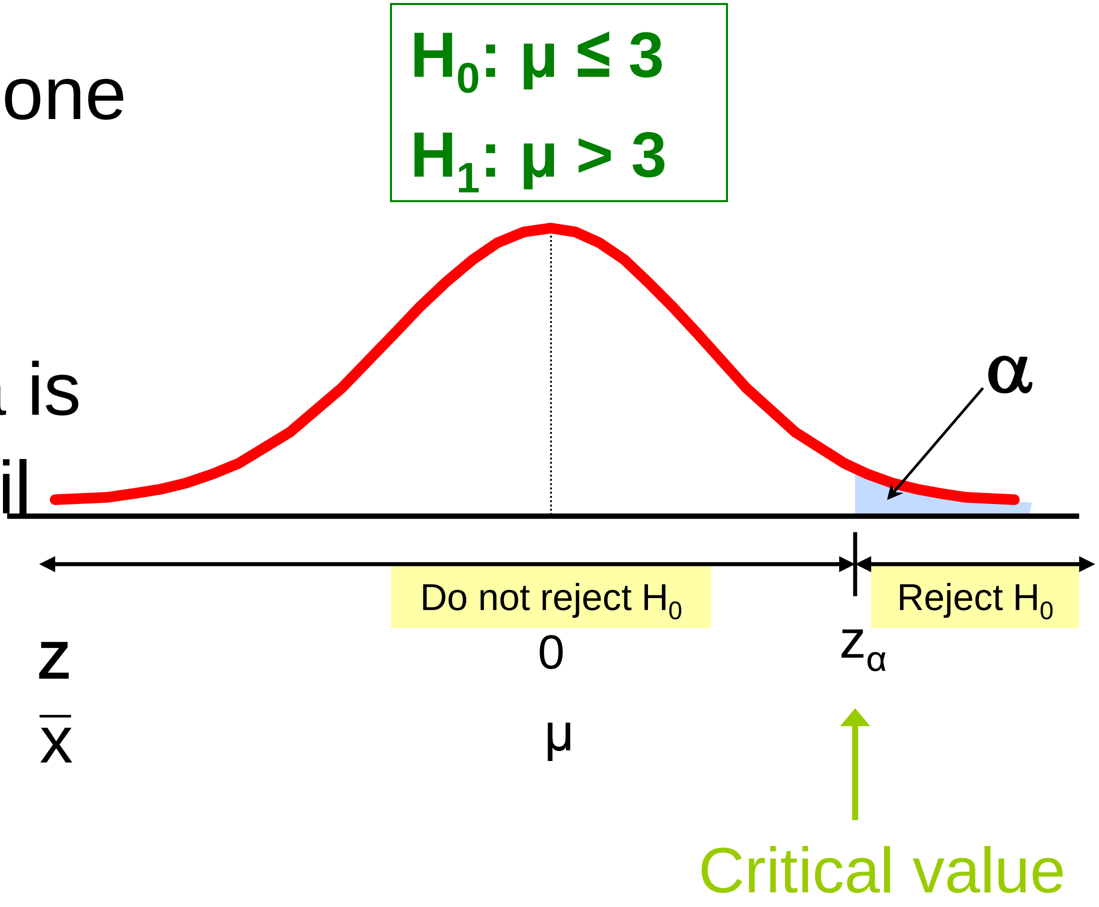

Lower-Tail Tests

- There is only one critical value, since the rejection area is in only one tail

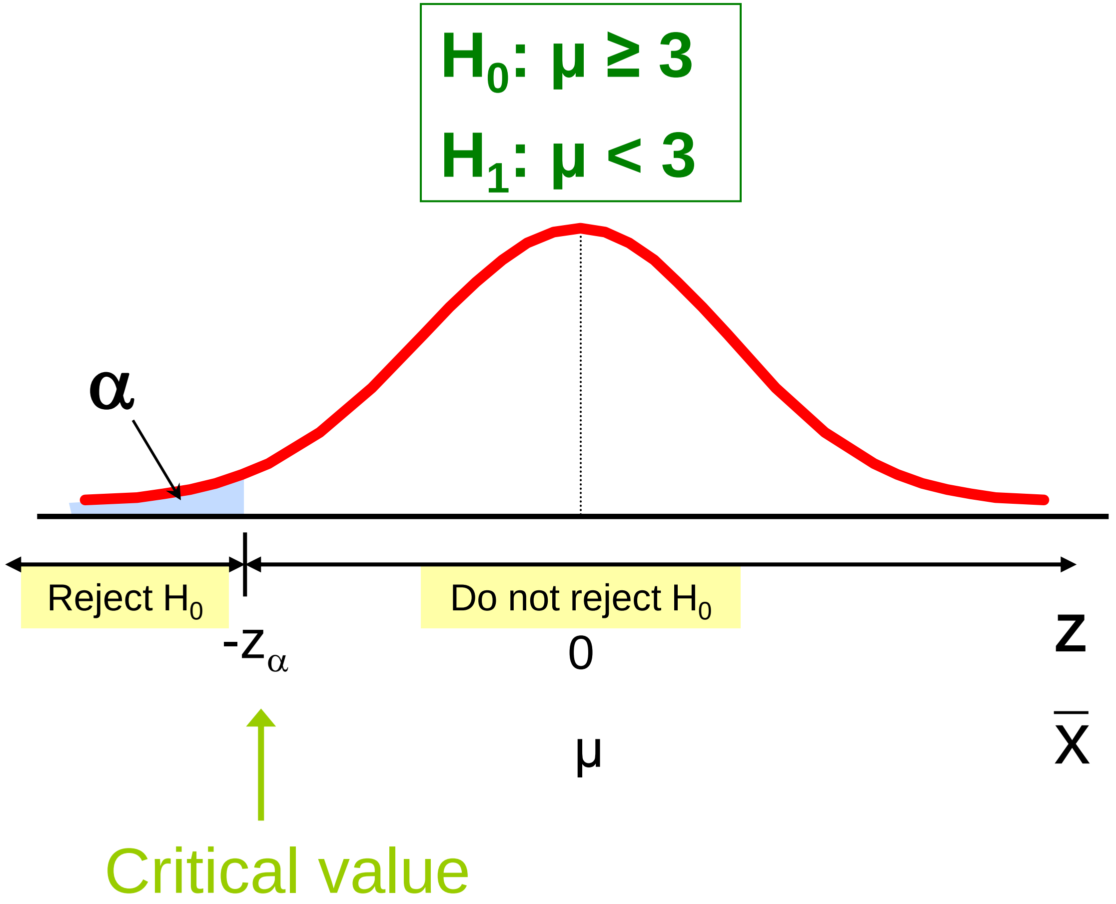

Two-Tail Tests

- In some settings, the alternative hypothesis does not specify a unique direction

H_0: $\mu = 3$
H_1: $\mu \neq 3$

- There are two critical values, defining the two regions of rejection

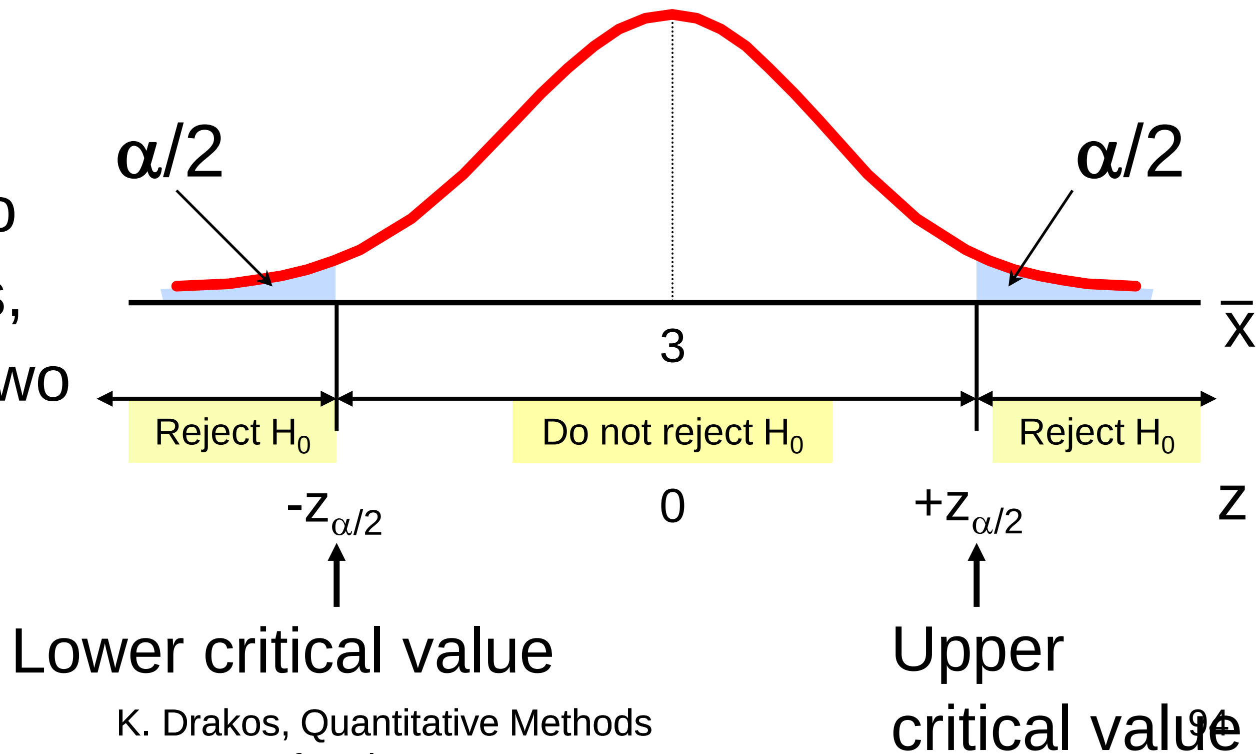

Hypothesis Testing Example

Test the claim that the true mean # of TV sets in US homes is equal to 3.
(Assume σ = 0.8)

- State the appropriate null and alternative hypotheses
  - H_0: μ = 3 , H_1: μ ≠ 3 (This is a two tailed test)
- Specify the desired level of significance
  - Suppose that α = .05 is chosen for this test
- Choose a sample size
  - Suppose a sample of size n = 100 is selected

Hypothesis Testing Example

(continued)

- Determine the appropriate technique
  - σ is known so this is a z test
- Set up the critical values
  - For α = .05 the critical z values are ±1.96
- Collect the data and compute the test statistic
  - Suppose the sample results are

    n = 100, $\bar{x}$ = 2.84 (σ = 0.8 is assumed known)

So the test statistic is:

$$z = \frac{\bar{X} - \mu_o}{\frac{\sigma}{\sqrt{n}}} = \frac{2.84 - 3}{\frac{0.8}{\sqrt{100}}} = \frac{-.16}{.08} = -2.0$$

Hypothesis Testing Example

(continued)

- Is the test statistic in the rejection region?

Reject H_0 if $z < -1.96$ or $z > 1.96$; otherwise do not reject H_0

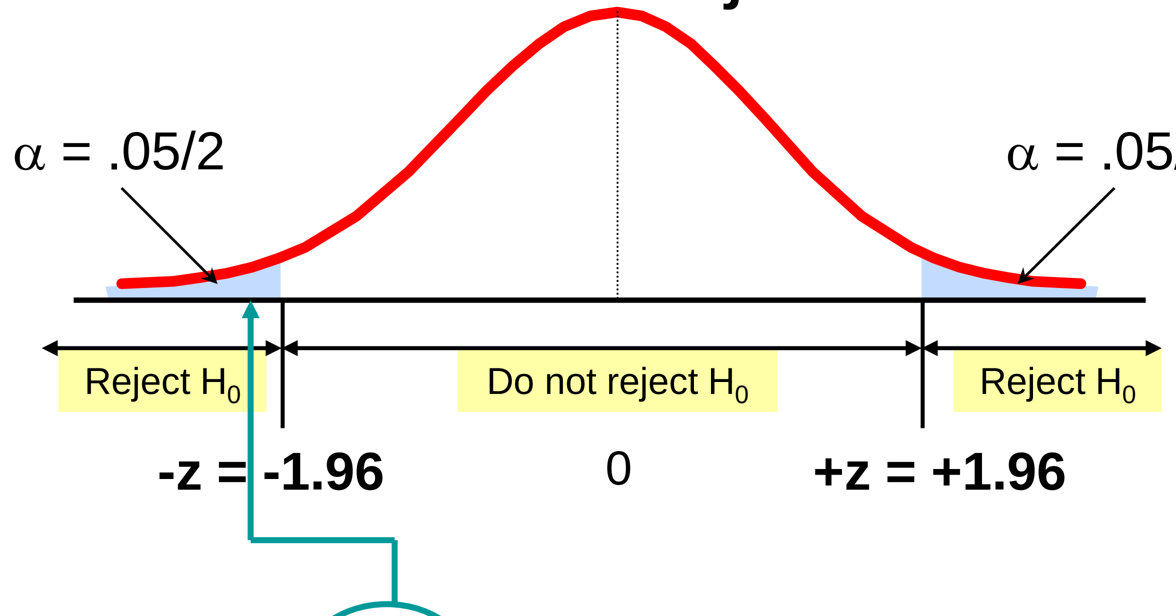

Here, $z = -2.0 < -1.96$, so the test statistic is in the rejection region

Hypothesis Testing Example

(continued)

- Reach a decision and interpret the result

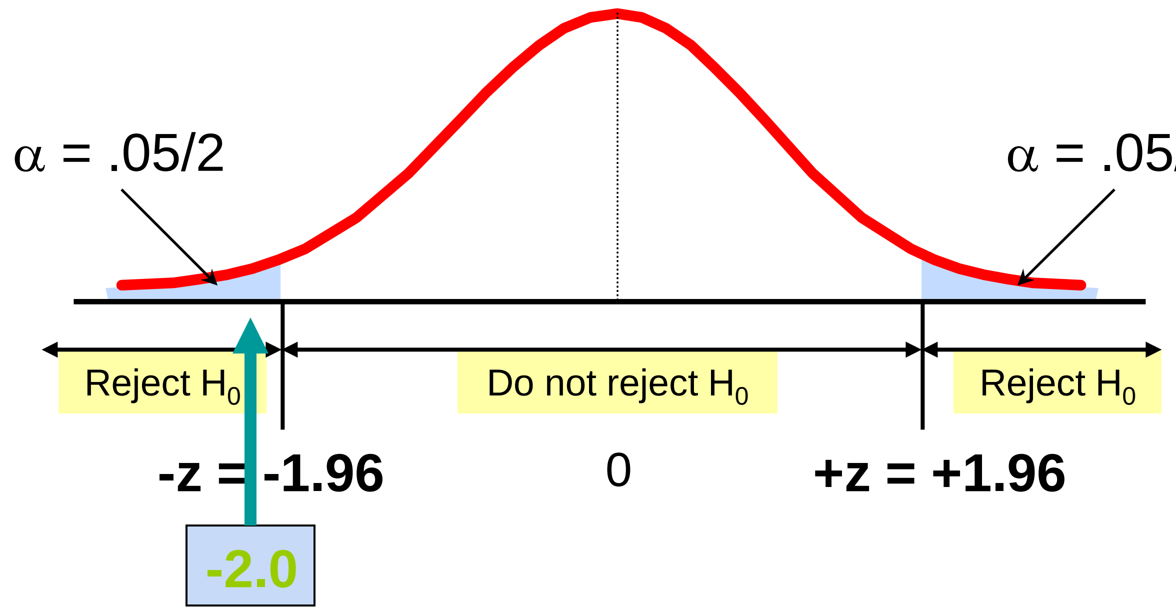

Since $z = -2.0 < -1.96$, we reject the null hypothesis and conclude that there is sufficient evidence that the mean number of TVs in US homes is not equal to 3

Example: p-Value

- Example: How likely is it to see a sample mean of 2.84 (or something further from the mean, in either direction) if the true mean is $\mu = 3.0$?

$\bar{x} = 2.84$ is translated to a z score of z = -2.0

$P(z < -2.0) = .0228$

$P(z > 2.0) = .0228$

p-value

= .0228 + .0228 = .0456

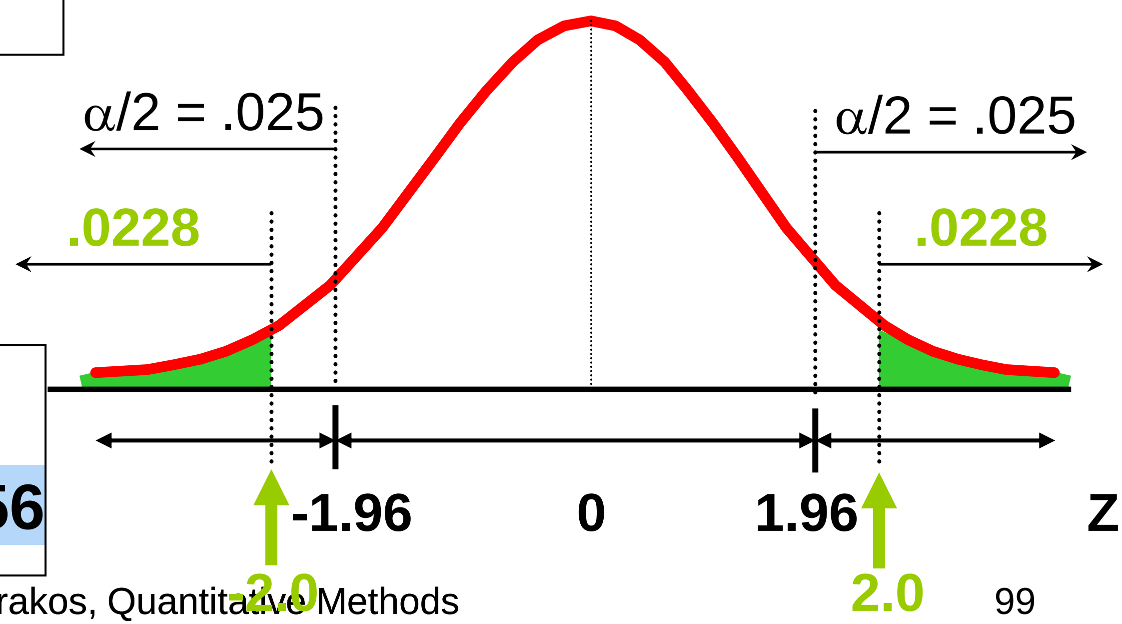

Example: p-Value

(continued)

- Compare the p-value with α
  - If p-value $<$ α , reject H_0
  - If p-value $\geq$ α , do not reject H_0

Here: p-value = .0456
α = .05

Since .0456 < .05, we reject the null hypothesis

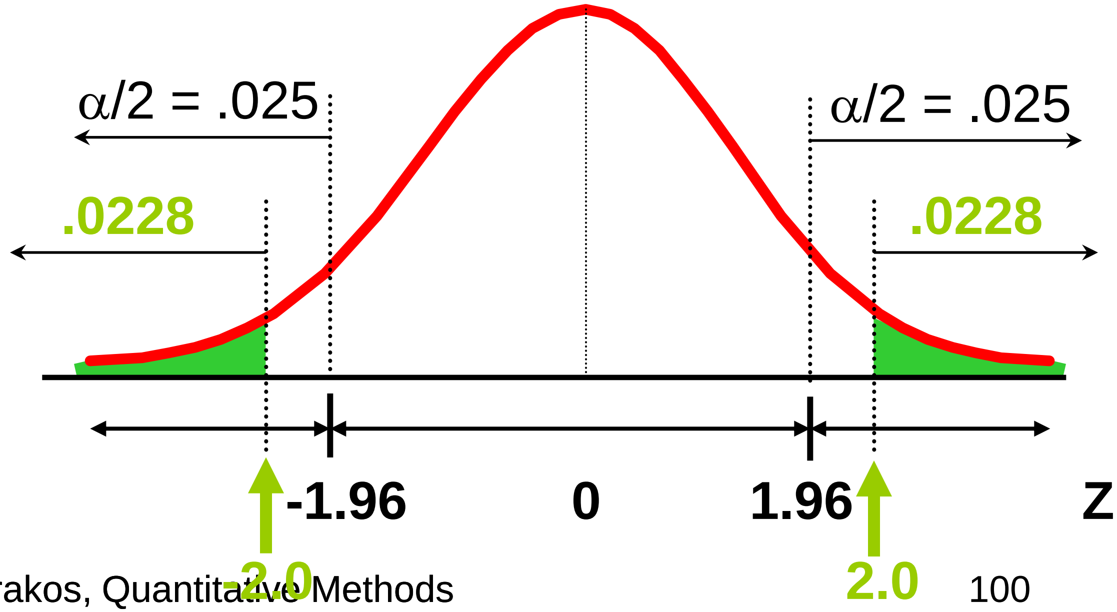

t Test of Hypothesis for the Mean (σ Unknown)

- Convert sample result ($\bar{x}$) to a t test statistic

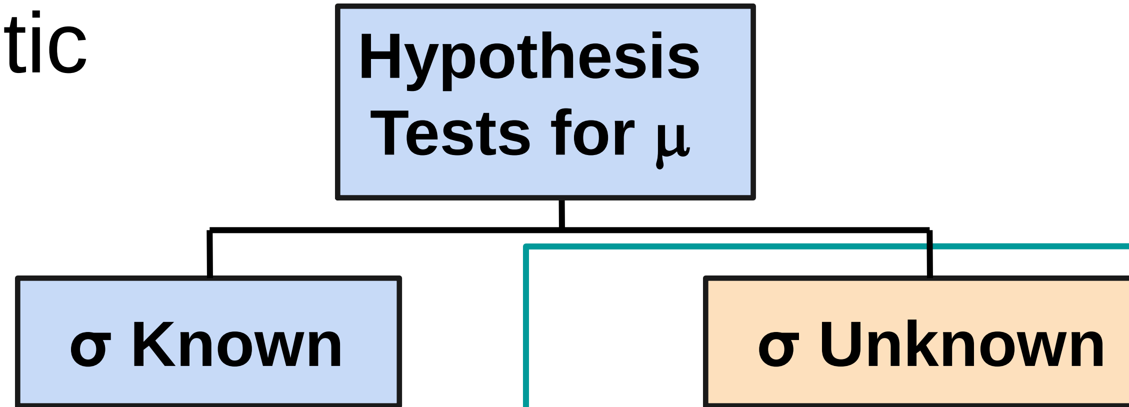

Consider the test

$$H_0 : \mu = \mu_0$$
$$H_1 : \mu > \mu_0$$

(Assume the population is normal)

The decision rule is:

$$\text{Reject } H_0 \text{ if } t = \frac{\bar{x} - \mu_0}{\frac{s}{\sqrt{n}}} > t_{n-1,\alpha}$$

t Test of Hypothesis for the Mean (σ Unknown)

(continued)

- For a two-tailed test:

Consider the test

$$H_0 : \mu = \mu_0$$
$$H_1 : \mu \neq \mu_0$$

(Assume the population is normal, and the population variance is unknown)

The decision rule is:

Reject H_0 if $t = \dfrac{\bar{x} - \mu_0}{\dfrac{s}{\sqrt{n}}} < -t_{n-1, \alpha/2}$ or if $t = \dfrac{\bar{x} - \mu_0}{\dfrac{s}{\sqrt{n}}} > t_{n-1, \alpha/2}$

Example: Two-Tail Test (σ Unknown)

The average cost of a 5-star hotel room in Athens is said to be 168 euros per night. A random sample of 25 hotels resulted in $\bar{x} = 172.50$ euros and

$s = 15.40$ euros. Test at the

$\alpha = 0.05$ level.

(Assume the population distribution is normal)

H_0: $\mu = 168$
H_1: $\mu \neq 168$

Example Solution: Two-Tail Test

$H_0: \mu = 168$
$H_1: \mu \neq 168$

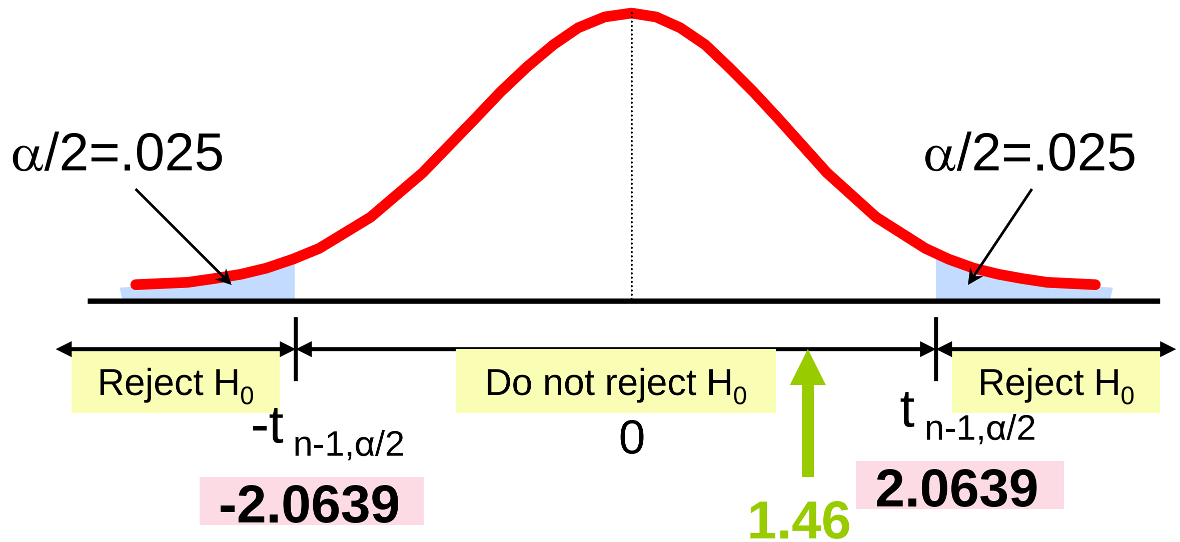

- n = 25
- a = 0.05
- σ is unknown, so use a t statistic
- Critical Value:

$t_{24\,,\,0.025} = \pm\, 2.0639$

$$t_{n-1} = \frac{\bar{x} - \mu}{\frac{s}{\sqrt{n}}} = \frac{172.50 - 168}{\frac{15.40}{\sqrt{25}}} = 1.46$$

Do not reject H_0: not sufficient evidence that true mean cost is different than \$168