SECTION 13.1 – Comparing Three or More Means (ANOVA)

Analysis of Variance

Analysis of Variance (ANOVA) is an inferential method used to test the equality of three or more population means.

Analysis-of-variance or ANOVA procedures rely on a distribution called the *F*-distribution, named in honor of Sir Ronald Fisher. A variable is said to have an *F*-distribution if its distribution has the shape of a special type of right-skewed curve, called an *F*-curve. There are infinitely many *F*-distributions, and we identify an *F*-distribution (and *F*-curve) by its number of degrees of freedom, just as we did for *t*-distributions and chi-square distributions. The difference, however, is that an *F*-distribution has two numbers of degrees of freedom instead of one.

The first number of degrees of freedom for an *F*-curve is called the degrees of freedom for the numerator. The second number of degrees of freedom for an *F*-curve is called the degrees of freedom for the denominator. Thus, the notation will look as follows:

df = (dfn, dfd) = (degrees of freedom for the numerator, degrees of freedom for the denominator)

Basic Properties of *F*-Curves

Property 1: The total area under a *F*-curve equals 1.

Property 2: A *F*-curve starts at 0 on the horizontal axis and extends indefinitely to the right, approaching, but never touching, the horizontal axis as it does so.

Property 3: A *F*-curve is right skewed.

Two different *F*-curves

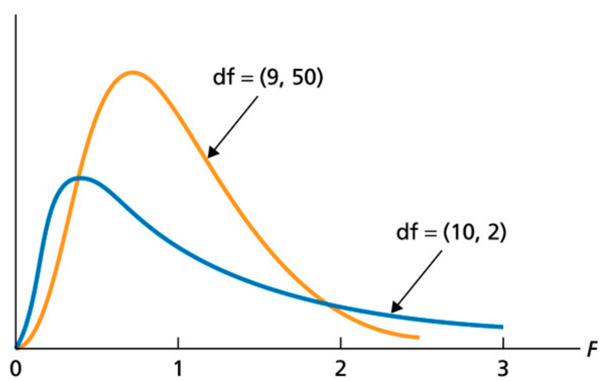

☺ **Exercises**:

Use Table IX to find the required F-values. Illustrate your work graphically.

1) An *F*-curve has df = (7,3).

a) find the *F*-value that has area 0.01 to its right.

b) find the *F*-value that has area 0.05 to its right.

2) For an *F*-curve with df = (6,10), find $F_{0.05}$ and $F_{0.025}$.

SECTION 13.1 – Comparing Three or More Means (ANOVA)

In Section 11.3 we used the non-pooled *t*-test to compare the means of two samples. When one assumes that the population standard deviations are equal, then a pooled *t*-test is required. If we have more than two means to compare, we could compare each mean with each other mean using many pooled *t*-tests. However, conducting multiple pooled *t*-tests increases the probability of Type I errors*. Instead Analysis of Variance (ANOVA) is used to compare two or more means of different populations.

Here we study the simplest kind of ANOVA called *one-way ANOVA*. One-way ANOVA compares the means of a variable for populations that result from classification by one other variable, known as the *factor*. The possible values of the factor are referred to as the *levels* of the factor.

One-way ANOVA is the generalization to more than two populations of the pooled *t*-test. (Note: Both one-way ANOVA and the pooled *t*-test yield the same result when applied to two populations.) In other words, *t*-tests are just a special case of ANOVA.

ANOVA is a type of hypothesis test like the pooled *t*-test. As such, it uses H_0 and H_1. When the hypotheses are used with three or more population means we determine the following:

1. All the means are the same (H_0); OR
2. At least one of the means is different from the other means (H_1). (But we don't know which one.)

* Every time you conduct a *t*-test there is a chance that you will make a Type I error. This error is usually 5%. By running two *t*-tests on the same data you will have increased your chance of "making a mistake" to 10%. The formula for determining the new error rate for multiple *t*-tests is not as simple as multiplying 5% by the number of tests. However, if you are only making a few multiple comparisons, the results are very similar if you do. As such, three *t*-tests would be 15% (actually, 14.3%) and so on. These are unacceptable errors. An ANOVA controls for these errors so that the Type I error remains at 5% and you can be more confident that any significant result you find is not just down to chance.

Mean Squares and *F*-Statistic in One-Way ANOVA

Treatment sum of squares, *SST*:

$$SST = n_1(\bar{x}_1 - \bar{x})^2 + n_2(\bar{x}_2 - \bar{x})^2 + \cdots + n_k(\bar{x}_k - \bar{x})^2$$

Treatment mean square, *MST*:

The variation among the sample means: $MST = \dfrac{SST}{(k-1)}$, where k is the number of populations under consideration.

Error sum of squares, *SSE*:

$$SSE = (n_1 - 1)s_1^{\,2} + (n_2 - 1)s_2^{\,2} + \cdots + (n_k - 1)s_k^{\,2}$$

Error mean square, *MSE*:

The variation within the samples: $MSE = \dfrac{SSE}{(n-k)}$, where n is the total number of observations.

***F*-Statistic, *F*:**

The ratio of the variation among the sample means to the variation within the samples:

$$F = \frac{\text{mean square due to treatment}}{\text{mean square due to error}} = \frac{MST}{MSE}$$

SECTION 13.1 – Comparing Three or More Means (ANOVA)

☺ **Example #1**:

Crash Data. The Insurance Institute for Highway Safety conducts experiments in which cars are crashed into a fixed barrier at 40 mph. In the institute's 40-mph offset test, 40% of the total width of each vehicle strikes a barrier on the driver's side. The barrier's deformable face is made of aluminum honeycomb, which makes the forces in the test similar to those involved in a frontal offset crash between two vehicles of the same weight, each going just less than 40 mph. You are in the market to buy a family car and you want to know if the mean head injury resulting from this offset crash is the same for large family cars, passenger vans, and midsize utility vehicles. The following data were collected from the institute's study.

Large Family Cars	Head Injury (hic)	Passenger Vans	Head Injury (hic)	Midsize Utility Vehicles	Head Injury (hic)
Hyundai XG300	264	Toyota Sienna	148	Honda Pilot	225
Ford Taurus	134	Honda Odyssey	238	Toyota 4Runner	216
Buick LaSabre	409	Ford Freestar	340	Mitsubishi Endeavor	186
Chevrolet Impala	530	Mazda MPV	693	Nissan Murano	307
Chrysler 300	149	Chevrolet Uplander	550	Ford Explorer	353
Pontiac Grand Prix	627	Nissan Quest	470	Kia Sorento	552
Toyota Avalon	166	Kia Sedona	322	Chevy Trailblazer	397

Determine each of the following, *SST*, *MST*, *SSE*, *MSE*, and *F*.

$$\begin{aligned} SST &= n_1(\bar{x}_1 - \bar{x})^2 + n_2(\bar{x}_2 - \bar{x})^2 + \cdots + n_k(\bar{x}_k - \bar{x})^2 \\ &= 7(325.5714286 - 346.4761905)^2 + 7(394.4285714 - 346.4761905)^2 + 7(319.4285714 - 346.4761905)^2 \\ &= 3059.063491 + 16096.01584 + 5121.015893 \\ &= 24276.09522 \end{aligned}$$

$$MST = \frac{SST}{(k-1)} = \frac{24276.09522}{3-1} = 12138.04761$$

$$\begin{aligned} SSE &= (n_1 - 1)s_1^2 + (n_2 - 1)s_2^2 + \cdots + (n_k - 1)s_k^2 \\ &= (7-1)(198.6377416)^2 + (7-1)(188.1770949)^2 + (7-1)(128.3262212)^2 \\ &= 236741.7143 + 212463.7143 + 98805.71428 \\ &= 548011.1429 \end{aligned}$$

$$MSE = \frac{SSE}{(n-k)} = \frac{548011.1429}{(21-3)} = 30445.06349$$

$$F = \frac{MST}{MSE} = \frac{12138.04761}{30445.06349} = 0.3986868878 \approx 0.399$$

SECTION 13.1 – Comparing Three or More Means (ANOVA)

Distribution of the *F*-Statistic for One-Way ANOVA
Suppose that the variable under consideration is normally distribution on each of the k populations and that the population standard deviations are equal. Then, for independent samples from the k populations, the variable

$$F = \frac{MST}{MSE}$$

has the F-distribution df = $(k-1, n-k)$ if the null hypothesis of equal population means is true. Here, n denotes the total number of observations.

HYPOTHESIS TEST – One-Way ANOVA Test
Purpose: To perform a hypothesis test to compare k population means $\mu_1, \mu_2, \ldots \mu_k$.
Assumptions:
1) Simple random samples **2)** Independent samples **3)** Normal populations
4) Equal population standard deviations

Step 1 – The null and alternative hypotheses are, respectively:

$H_0 : \mu_1 = \mu_2 = \cdots = \mu_k$
H_1 : At least one mean is different.

Step 2 – Decide on the significance level α.

Step 3 – Compute the value of the test statistic $F = \dfrac{MST}{MSE}$, and denote that value F_0.

CLASSICAL APPROACH OR ***P*-VALUE APPROACH**

CLASSICAL APPROACH

The critical value is F_α with df = $(k-1, n-k)$. Use Table IX to find the critical value.

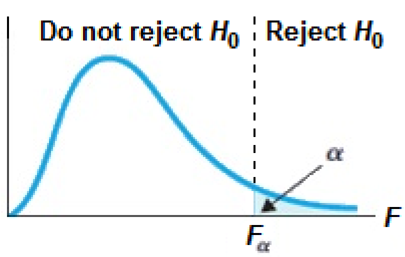

Step 4 – If the value of the test statistic falls in the rejection region, reject H_0; otherwise, do not reject H_0.

***P*-VALUE APPROACH**

The F-statistic has df = $(k-1, n-k)$. Use Table IX to estimate the P-value, or obtain it exactly using technology.

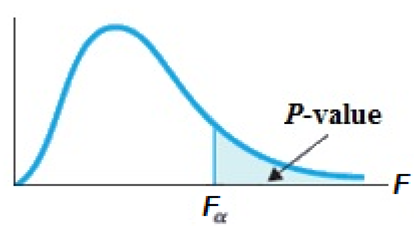

Step 4 – If $P \leq \alpha$, reject H_0; otherwise, do not reject H_0.

Step 5 – Interpret the results of the hypothesis test.

SECTION 13.1 – Comparing Three or More Means (ANOVA)

☺ **Exercises**:

3) **Crash Data.** The Insurance Institute for Highway Safety conducts experiments in which cars are crashed into a fixed barrier at 40 mph. In the institute's 40-mph offset test, 40% of the total width of each vehicle strikes a barrier on the driver's side. The barrier's deformable face is made of aluminum honeycomb, which makes the forces in the test similar to those involved in a frontal offset crash between two vehicles of the same weight, each going just less than 40 mph. You are in the market to buy a family car and you want to know if the mean head injury resulting from this offset crash is the same for large family cars, passenger vans, and midsize utility vehicles. The following data were collected from the institute's study.

Large Family Cars	**Head Injury (hic)**	**Passenger Vans**	**Head Injury (hic)**	**Midsize Utility Vehicles**	**Head Injury (hic)**
Hyundai XG300	264	Toyota Sienna	148	Honda Pilot	225
Ford Taurus	134	Honda Odyssey	238	Toyota 4Runner	216
Buick LaSabre	409	Ford Freestar	340	Mitsubishi Endeavor	186
Chevrolet Impala	530	Mazda MPV	693	Nissan Murano	307
Chrysler 300	149	Chevrolet Uplander	550	Ford Explorer	353
Pontiac Grand Prix	627	Nissan Quest	470	Kia Sorento	552
Toyota Avalon	166	Kia Sedona	322	Chevy Trailblazer	397

Using the hypothetical data provided above to test whether the mean head injury for each vehicle type is the same at the $\alpha = 0.01$ level of significance.

(Note: Recall from Example #3, that the test statistic calculated was $F = 0.399$).