Lecture 17:
Chapters 11-12

C C Moxley

UAB Mathematics

27 July 15

§11.1 Analyzing Categorical Data

So far, we have used statistical methods to analyze population parameters - but we knew what types of random variables we were dealing with.

§11.1 Analyzing Categorical Data

So far, we have used statistical methods to analyze population parameters - but we knew what types of random variables we were dealing with. Now, we move on to making qualitative inferences, namely things like whether or not a sample comes from a particular type of distribution or whether or not two samples are independent.

So far, we have used statistical methods to analyze population parameters - but we knew what types of random variables we were dealing with. Now, we move on to making qualitative inferences, namely things like whether or not a sample comes from a particular type of distribution or whether or not two samples are independent. The first type of test is a **goodness-of-fit test**,

§11.1 Analyzing Categorical Data

So far, we have used statistical methods to analyze population parameters - but we knew what types of random variables we were dealing with. Now, we move on to making qualitative inferences, namely things like whether or not a sample comes from a particular type of distribution or whether or not two samples are independent. The first type of test is a **goodness-of-fit test**, and the second test is done using a is a **contingency test**.

§11.2 Goodness-of-Fit Tests

Definition (Goodness-of-Fit Test)

A **goodness-of-fit test** tests the hypothesis that an observed frequency distribution or a single column of data comes from some claimed distribution.

§11.2 Goodness-of-Fit Tests

Definition (Goodness-of-Fit Test)

A **goodness-of-fit test** tests the hypothesis that an observed frequency distribution or a single column of data comes from some claimed distribution.

There are a few requirements to perform a goodness-of-fit test.

§11.2 Goodness-of-Fit Tests

Definition (Goodness-of-Fit Test)

A **goodness-of-fit test** tests the hypothesis that an observed frequency distribution or a single column of data comes from some claimed distribution.

There are a few requirements to perform a goodness-of-fit test.

- The data must be randomly selected.

§11.2 Goodness-of-Fit Tests

Definition (Goodness-of-Fit Test)

A **goodness-of-fit test** tests the hypothesis that an observed frequency distribution or a single column of data comes from some claimed distribution.

There are a few requirements to perform a goodness-of-fit test.

- The data must be randomly selected.
- The data consist of (or can be arranged into) frequency counts for each of the different categories.

§11.2 Goodness-of-Fit Tests

Definition (Goodness-of-Fit Test)

A **goodness-of-fit test** tests the hypothesis that an observed frequency distribution or a single column of data comes from some claimed distribution.

There are a few requirements to perform a goodness-of-fit test.

- The data must be randomly selected.
- The data consist of (or can be arranged into) frequency counts for each of the different categories.
- For each category, the expected frequency is at least five.

We need these variables.

§11.2 Goodness-of-Fit Tests

We need these variables.

- O, observed frequency (from the table or data)

§11.2 Goodness-of-Fit Tests

We need these variables.

- O, observed frequency (from the table or data)
- E, expected frequency (from the assumption)

We need these variables.

- O, observed frequency (from the table or data)
- E, expected frequency (from the assumption)
- k, number of different cells or categories

§11.2 Goodness-of-Fit Tests

We need these variables.

- O, observed frequency (from the table or data)
- E, expected frequency (from the assumption)
- k, number of different cells or categories
- n, number of trails or observations

§11.2 Goodness-of-Fit Tests

We need these variables.

- O, observed frequency (from the table or data)
- E, expected frequency (from the assumption)
- k, number of different cells or categories
- n, number of trails or observations

We then compute the test statistic $\chi^2 = \sum \frac{(O-E)^2}{E}$.

§11.2 Goodness-of-Fit Tests

We need these variables.

- O, observed frequency (from the table or data)
- E, expected frequency (from the assumption)
- k, number of different cells or categories
- n, number of trails or observations

We then compute the test statistic $\chi^2 = \sum \frac{(O-E)^2}{E}$. And we perform a right-tailed χ^2-test with $k-1$ degrees of freedom.

How to calculate E?

How to calculate E? Well, if the assumed distribution is uniform, then we get that $E = n/k$. Otherwise, we get $E = np$, where p is the probability of falling into a particular category.

How to calculate E? Well, if the assumed distribution is uniform, then we get that $E = n/k$. Otherwise, we get $E = np$, where p is the probability of falling into a particular category. You can often calculate p by using a "between" calculator on StatCrunch for the appropriate distribution.

§11.2 Example

A salesman wants to make a pitch on a day a company is most likely to make major purchases. He obtained data about the company's previous purchases and tested (using $\alpha = 0.05$) to see if the purchases were uniformly distributed throughout the business week (Mon-Fri). He obtained the following.

- $k = 7$,

§11.2 Example

A salesman wants to make a pitch on a day a company is most likely to make major purchases. He obtained data about the company's previous purchases and tested (using $\alpha = 0.05$) to see if the purchases were uniformly distributed throughout the business week (Mon-Fri). He obtained the following.

- $k = 7$, $n = 20$,

§11.2 Example

A salesman wants to make a pitch on a day a company is most likely to make major purchases. He obtained data about the company's previous purchases and tested (using $\alpha = 0.05$) to see if the purchases were uniformly distributed throughout the business week (Mon-Fri). He obtained the following.

- $k = 7$, $n = 20$, $E = 2.86$,

§11.2 Example

A salesman wants to make a pitch on a day a company is most likely to make major purchases. He obtained data about the company's previous purchases and tested (using $\alpha = 0.05$) to see if the purchases were uniformly distributed throughout the business week (Mon-Fri). He obtained the following.

- $k = 7$, $n = 20$, $E = 2.86$, test statistic $\chi^2 = 15.85$, critical value 15.507, P-value 0.045

What's the null hypothesis in this case?

§11.2 Example

A salesman wants to make a pitch on a day a company is most likely to make major purchases. He obtained data about the company's previous purchases and tested (using $\alpha = 0.05$) to see if the purchases were uniformly distributed throughout the business week (Mon-Fri). He obtained the following.

- $k = 7$, $n = 20$, $E = 2.86$, test statistic $\chi^2 = 15.85$, critical value 15.507, P-value 0.045

What's the null hypothesis in this case? Do we accept or reject it?

§11.2 Example

A salesman wants to make a pitch on a day a company is most likely to make major purchases. He obtained data about the company's previous purchases and tested (using $\alpha = 0.05$) to see if the purchases were uniformly distributed throughout the business week (Mon-Fri). He obtained the following.

- $k = 7$, $n = 20$, $E = 2.86$, test statistic $\chi^2 = 15.85$, critical value 15.507, P-value 0.045

What's the null hypothesis in this case? Do we accept or reject it? What does this mean?

§11.2 Example (with frequency table)

A casino wants to determine if a 10-sided die is loaded (with significance $\alpha = 0.01$). They roll the die 64 times and get the following data:

§11.2 Example (with frequency table)

A casino wants to determine if a 10-sided die is loaded (with significance $\alpha = 0.01$). They roll the die 64 times and get the following data:

Side	1	2	3	4	5	6	7	8	9	10
Rolls	6	6	5	7	7	7	7	7	5	7

§11.2 Example (with frequency table)

A casino wants to determine if a 10-sided die is loaded (with significance $\alpha = 0.01$). They roll the die 64 times and get the following data:

Side	1	2	3	4	5	6	7	8	9	10
Rolls	6	6	5	7	7	7	7	7	5	7

We end up getting that the P-value for this test is 0.9994, so we **fail to reject** the null hypothesis.

§11.2 Example (with frequency table)

A casino wants to determine if a 10-sided die is loaded (with significance $\alpha = 0.01$). They roll the die 64 times and get the following data:

Side	1	2	3	4	5	6	7	8	9	10
Rolls	6	6	5	7	7	7	7	7	5	7

We end up getting that the P-value for this test is 0.9994, so we **fail to reject** the null hypothesis. This means that the data does not support the claim that the die is loaded.

§11.2 Example (with frequency table)

A casino wants to determine if a 10-sided die is loaded (with significance $\alpha = 0.01$). They roll the die 64 times and get the following data:

Side	1	2	3	4	5	6	7	8	9	10
Rolls	6	6	5	7	7	7	7	7	5	7

We end up getting that the P-value for this test is 0.9994, so we **fail to reject** the null hypothesis. This means that the data does not support the claim that the die is loaded. We can conduct the same test with a critical value. We'd see that our test statistic was 1 and our critical value was 21.666, telling us that we should reject our null hypothesis as well.

§11.2 Example (with frequency table)

A casino wants to determine if a 10-sided die is loaded (with significance $\alpha = 0.01$). They roll the die 64 times and get the following data:

Side	1	2	3	4	5	6	7	8	9	10
Rolls	6	6	5	7	7	7	7	7	5	7

We end up getting that the P-value for this test is 0.9994, so we **fail to reject** the null hypothesis. This means that the data does not support the claim that the die is loaded. We can conduct the same test with a critical value. We'd see that our test statistic was 1 and our critical value was 21.666, telling us that we should reject our null hypothesis as well. Use the Goodness-of-Fit tab in StatCrunch.

§11.3 Contingency (Two-Way Frequency) Tables

Definition (Contingency Table)

A **contingency table** is a table consisting of frequency counts of categorical data corresponding to two different variables.

§11.3 Contingency (Two-Way Frequency) Tables

Definition (Contingency Table)

A **contingency table** is a table consisting of frequency counts of categorical data corresponding to two different variables.

Definition (Test of Independence)

A **test of independence** tests the claim that the row and column variables are independent.

§11.3 Contingency (Two-Way Frequency) Tables

Definition (Contingency Table)

A **contingency table** is a table consisting of frequency counts of categorical data corresponding to two different variables.

Definition (Test of Independence)

A **test of independence** tests the claim that the row and column variables are independent.

A test of independence is a χ^2 test, and its degrees of freedom is computed by finding the product of one fewer than the number of columns and one fewer than the number of rows in the contingency table: $df = (c-1)(r-1)$.

§11.3 Contingency (Two-Way Frequency) Tables

The requirements for a test of independence are below.

- The sample data must be randomly selected.

§11.3 Contingency (Two-Way Frequency) Tables

The requirements for a test of independence are below.

- The sample data must be randomly selected.
- The sample data are frequency counts in a two-way table.

§11.3 Contingency (Two-Way Frequency) Tables

The requirements for a test of independence are below.

- The sample data must be randomly selected.
- The sample data are frequency counts in a two-way table.
- Each cell in the **expected** table has a value at least 5. (The expected table is the table we'd get if the columns and rows were independent.)

The requirements for a test of independence are below.

- The sample data must be randomly selected.
- The sample data are frequency counts in a two-way table.
- Each cell in the **expected** table has a value at least 5. (The expected table is the table we'd get if the columns and rows were independent.)

You can read more details about the test of independence on page 578 of your text.

§11.3 Example

A hospital wants to verify that a test for diabetes is effective by seeing if the result of the test and the status of the patient are dependent with significance $\alpha = 0.01$. They have the following data.

	Negative	Positive
Tested Negative	17	7
Tested Positive	8	15

§11.3 Example

A hospital wants to verify that a test for diabetes is effective by seeing if the result of the test and the status of the patient are dependent with significance $\alpha = 0.01$. They have the following data.

	Negative	Positive
Tested Negative	17	7
Tested Positive	8	15

Use the contingency table with summary tab in StatCrunch to see that the P-value is 0.0133. So, we fail to reject the null hypothesis! This means that the data supports the claim that the result of the test is independent of the status of the patient, meaning that the diabetes test is not effective.

§11.3 Example

Note: We would reach the same conclusion if we used a critical value test because the test is right-tailed and the critical value would be 6.635 with a test statistic of 6.131.

§12.1 ANOVA (Analysis of Variance) Tests

In Chapter 9, we tested to see if two population means were equal.

§12.1 ANOVA (Analysis of Variance) Tests

In Chapter 9, we tested to see if two population means were equal. In this chapter, we will test to see if three or more population samples are the same using a **one-way** ANOVA test. We call it "one-way" because we separate our populations into groups based on one characteristic.

§12.1 ANOVA (Analysis of Variance) Tests

In Chapter 9, we tested to see if two population means were equal. In this chapter, we will test to see if three or more population samples are the same using a **one-way** ANOVA test. We call it "one-way" because we separate our populations into groups based on one characteristic. This is often called an "analysis of variance" but this refers to the method of testing, not the thing which we are testing - which is means not variances.

§12.1 ANOVA (Analysis of Variance) Tests

In Chapter 9, we tested to see if two population means were equal. In this chapter, we will test to see if three or more population samples are the same using a **one-way** ANOVA test. We call it "one-way" because we separate our populations into groups based on one characteristic. This is often called an "analysis of variance" but this refers to the method of testing, not the thing which we are testing - which is means not variances. We will also compare the population means of populations separated into categories based on two characteristics (such as weight and body temperature).

§12.2 One-Way ANOVA Test

Definition (One-Way ANOVA Test)

A one-way ANOVA test is a method of testing the equality of three or more population means by analyzing sample variances. One-way analysis of variance is used with data categorized with one **factor**, so there is one characteristic used to separate the sample data into the different categories.

§12.2 One-Way ANOVA Test

Definition (One-Way ANOVA Test)

A one-way ANOVA test is a method of testing the equality of three or more population means by analyzing sample variances. One-way analysis of variance is used with data categorized with one **factor**, so there is one characteristic used to separate the sample data into the different categories.

- Populations must be approximately normally distributed.

§12.2 One-Way ANOVA Test

Definition (One-Way ANOVA Test)

A one-way ANOVA test is a method of testing the equality of three or more population means by analyzing sample variances. One-way analysis of variance is used with data categorized with one **factor**, so there is one characteristic used to separate the sample data into the different categories.

- Populations must be approximately normally distributed.
- They must have the same variance, roughly. (Or at least the sample sizes must be the same for each category.)

§12.2 One-Way ANOVA Test

Definition (One-Way ANOVA Test)

A one-way ANOVA test is a method of testing the equality of three or more population means by analyzing sample variances. One-way analysis of variance is used with data categorized with one **factor**, so there is one characteristic used to separate the sample data into the different categories.

- Populations must be approximately normally distributed.
- They must have the same variance, roughly. (Or at least the sample sizes must be the same for each category.)
- The samples must be simple random and quantitative.

§12.2 One-Way ANOVA Test

Definition (One-Way ANOVA Test)

A one-way ANOVA test is a method of testing the equality of three or more population means by analyzing sample variances. One-way analysis of variance is used with data categorized with one **factor**, so there is one characteristic used to separate the sample data into the different categories.

- Populations must be approximately normally distributed.
- They must have the same variance, roughly. (Or at least the sample sizes must be the same for each category.)
- The samples must be simple random and quantitative.
- The samples must be independent - not matched or paired.

§12.2 One-Way ANOVA Test

Definition (One-Way ANOVA Test)

A one-way ANOVA test is a method of testing the equality of three or more population means by analyzing sample variances. One-way analysis of variance is used with data categorized with one **factor**, so there is one characteristic used to separate the sample data into the different categories.

- Populations must be approximately normally distributed.
- They must have the same variance, roughly. (Or at least the sample sizes must be the same for each category.)
- The samples must be simple random and quantitative.
- The samples must be independent - not matched or paired.
- The different samples are from populations that are categorized in only one way.

§12.2 One-Way ANOVA Test

In a one-way ANOVA test, our null hypothesis is always that all the population means are equal. The alternative hypothesis, then, is that at least one is different from the others.

§12.2 One-Way ANOVA Test

In a one-way ANOVA test, our null hypothesis is always that all the population means are equal. The alternative hypothesis, then, is that at least one is different from the others.

$$H_0 : \mu_1 = \mu_2 = \cdots = \mu_n$$

§12.2 One-Way ANOVA Test

In a one-way ANOVA test, our null hypothesis is always that all the population means are equal. The alternative hypothesis, then, is that at least one is different from the others.

$$H_0 : \mu_1 = \mu_2 = \cdots = \mu_n$$

H_1 : At least one of the population means differs from the others.

§12.2 One-Way ANOVA Test

In a one-way ANOVA test, our null hypothesis is always that all the population means are equal. The alternative hypothesis, then, is that at least one is different from the others.

$$H_0 : \mu_1 = \mu_2 = \cdots = \mu_n$$

H_1 : At least one of the population means differs from the others.

You can find "by hand" computations for a one-way ANOVA test, see page 605. The test is right-tailed.

§12.2 Example

The weights (in kg) of oak trees is given below for trees planted in the same plot but with different growth-enhancing methods applied. Use a 0.05 significance level to test the claim that the four treatment categories yield oak trees with the same mean weight. Does there appear to be a best method for enhancing growth of trees in this soil?

Control	Fert	Irr	Fert & Irr
0.24	0.92	0.96	1.07
1.69	0.07	1.43	1.63
1.23	0.56	1.26	1.39
0.99	1.74	1.57	0.49
0.99	1.74	1.57	0.49

§12.2 Example

Control	Fert	Irr	Fert & Irr
0.24	0.92	0.96	1.07
1.69	0.07	1.43	1.63
1.23	0.56	1.26	1.39
0.99	1.74	1.57	0.49
1.80	1.13	0.72	0.95

We use StatCrunch to conduct the test.

§12.2 Example

What's our null and alternative hypotheses? Based on the output, do we reject or accept the null hypothesis? What does this mean? (Use $\alpha = 0.05$.)

Options

Analysis of Variance results:
Data stored in separate columns.

Column statistics

Column	n	Mean	Std. Dev.	Std. Error
Control	5	1.19	0.62573956	0.27983924
Fert	5	0.884	0.624924	0.27947451
Irr	5	1.188	0.34665545	0.15502903
F & I	5	1.106	0.43598165	0.19497692

ANOVA table

Source	DF	SS	MS	F-Stat	P-value
Columns	3	0.3114	0.1038	0.38010491	0.7687
Error	16	4.36932	0.2730825		
Total	19	4.68072			

§12.2 Example

You can also perform the previous test using a critical value test. The F critical value is computed using the significance, the F calculator, and the "degrees of freedom of the numerator" and the "degrees of freedom of the denominator" which are, respectively, one fewer than the number of columns and the difference between the total number of data points and the number of columns.

§12.2 Example

You can also perform the previous test using a critical value test. The F critical value is computed using the significance, the F calculator, and the "degrees of freedom of the numerator" and the "degrees of freedom of the denominator" which are, respectively, one fewer than the number of columns and the difference between the total number of data points and the number of columns. Recall the F stat from the output was 0.3801.

§12.2 Example

You can also perform the previous test using a critical value test. The F critical value is computed using the significance, the F calculator, and the "degrees of freedom of the numerator" and the "degrees of freedom of the denominator" which are, respectively, one fewer than the number of columns and the difference between the total number of data points and the number of columns. Recall the F stat from the output was 0.3801.

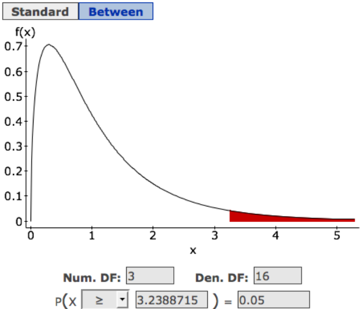

§12.3 Two-Way ANOVA Test

Definition (Two-Way ANOVA Test)

A two-way ANOVA test is a method of testing the effect of interactions of two factors and of testing (if the two factors do not interact) the effect of the factors separately.

§12.3 Two-Way ANOVA Test

Definition (Two-Way ANOVA Test)

A two-way ANOVA test is a method of testing the effect of interactions of two factors and of testing (if the two factors do not interact) the effect of the factors separately.

- Populations must be approximately normally distributed.

§12.3 Two-Way ANOVA Test

Definition (Two-Way ANOVA Test)

A two-way ANOVA test is a method of testing the effect of interactions of two factors and of testing (if the two factors do not interact) the effect of the factors separately.

- Populations must be approximately normally distributed.
- They must have the same variance, roughly.

§12.3 Two-Way ANOVA Test

Definition (Two-Way ANOVA Test)

A two-way ANOVA test is a method of testing the effect of interactions of two factors and of testing (if the two factors do not interact) the effect of the factors separately.

- Populations must be approximately normally distributed.
- They must have the same variance, roughly.
- The samples must be simple random and quantitative.

§12.3 Two-Way ANOVA Test

Definition (Two-Way ANOVA Test)

A two-way ANOVA test is a method of testing the effect of interactions of two factors and of testing (if the two factors do not interact) the effect of the factors separately.

- Populations must be approximately normally distributed.
- They must have the same variance, roughly.
- The samples must be simple random and quantitative.
- The samples must be independent - not matched or paired.

§12.3 Two-Way ANOVA Test

Definition (Two-Way ANOVA Test)

A two-way ANOVA test is a method of testing the effect of interactions of two factors and of testing (if the two factors do not interact) the effect of the factors separately.

- Populations must be approximately normally distributed.
- They must have the same variance, roughly.
- The samples must be simple random and quantitative.
- The samples must be independent - not matched or paired.
- The different samples are from populations that are categorized in two ways.

Definition (Two-Way ANOVA Test)

A two-way ANOVA test is a method of testing the effect of interactions of two factors and of testing (if the two factors do not interact) the effect of the factors separately.

- Populations must be approximately normally distributed.
- They must have the same variance, roughly.
- The samples must be simple random and quantitative.
- The samples must be independent - not matched or paired.
- The different samples are from populations that are categorized in two ways.
- All cells have the same number of values.

§12.3 Two-Way ANOVA Test

The null hypothesis is always that there is **no effect due to interaction between the two factors**. The alternative is that there is an effect from the interaction between the two factors. The test statistic is always

$$F = \frac{MS(\text{interaction})}{MS(\text{error})}.$$

Note: If you conclude that the there is an effect due to interaction, then you would not investigate the rows and columns separately.

§12.3 Example

Below are pulse rates of men and women over and under the age of 30. Conduct a two-way ANOVA and state the results. Use a 0.05 significance level.

	< 30	≥ 30
F	78 104 78 64 60 98 82 98 90 96	76 76 72 66 72 78 62 72 74 56
M	60 80 56 68 68 74 74 68 62 56	46 70 62 66 90 80 60 58 64 60

§12.3 Example

We get the following result from StatCrunch.

Analysis of Variance results:
Responses: Results
Row factor: Row Factor
Column factor: Column Factor

ANOVA table

Source	DF	SS	MS	F-Stat	P-value
Row Factor	1	1322.5	1322.5	10.862919	0.0022
Column Factor	1	592.9	592.9	4.8700374	0.0338
Interaction	1	448.9	448.9	3.6872319	0.0628
Error	36	4382.8	121.74444		
Total	39	6747.1			

§12.3 Example

So there it does not appear that there is an effect due to the interaction of age and gender, but there does seem to be an effect due to age and a separate effect due to gender.