Tuesday, April 23, 2019

- **Warm-up**
  - **Reese's advertises that their color distribution is 50% orange. When I counted out the colors in my bag, only 123 out of 307 were orange. Is this significant evidence that their advertised distribution is incorrect?**

- **Reese's Pieces Lab**

Objectives

- Content Objective: I will perform a X^2 test on data and describe clearly what I am testing.
- Social Objective: I will participate in the class activities with my group.
- Language Objective: I will clearly explain verbally what a X^2 test is and how it is different from the past tests.

Warm-up

Reese's advertises that their color distribution is 50% orange. When I counted out the colors in my bag, only 123 out of 307 were orange. Is this significant evidence that their advertised distribution is incorrect?

$H_0: p = 0.5$

$H_A: p \neq 0.5$

$z = -3.48$

P-val = 0.00049

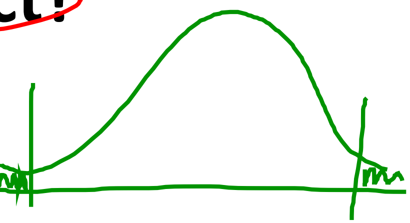

Representative Sample

307 < 10% of all R.P.

(.5)(307) ≥ 10 } success
(.5)(307) ≥ 10 } Fail

1 prop Z-test

Due to a p-value of 0.00049, which is lower than my α of 0.05, we reject the null. There is sufficient statistical evidence that their advertised distribution is incorrect.

Brown -> 38 + 35 + 24 + 44 + 36 + 47 + 42 + 32 = 298

Orange -> 66 + 48 + 54 + 57 + 62 + 86 + 72 + 58 = 503

Yellow -> 33 + 29 + 19 + 52 + 36 + 55 + 36 + 33 = 293

Counting Reese's Pieces

χ^2 Goodness-of-Fit

WHEN DO WE USE THIS INSTEAD OF Z OR T?

A test of whether the distribution of counts in one categorical variable matches the distribution predicted by a model is called a goodness-of-fit test.

As usual, we have 4 steps…

Running the test

- **How do I compare our count to the company?**
  - **Hypothesis** H_0: Population distribution matches advertised.
    H_A: Population distribution does not match advertised.
  - **Conditions**
  - **Mechanics**
  - **Conclusion**

Assumptions and Conditions

<u>Counted Data Condition</u>

Check that the data are ***<u>counts</u>*** for the categories of a categorical variable.

Assumptions and Conditions

- **Independence Assumption:** The counts in the cells should be independent of each other.
  - **Randomization Condition:** The individuals who have been counted and whose counts are available for analysis should be a random sample from some population.

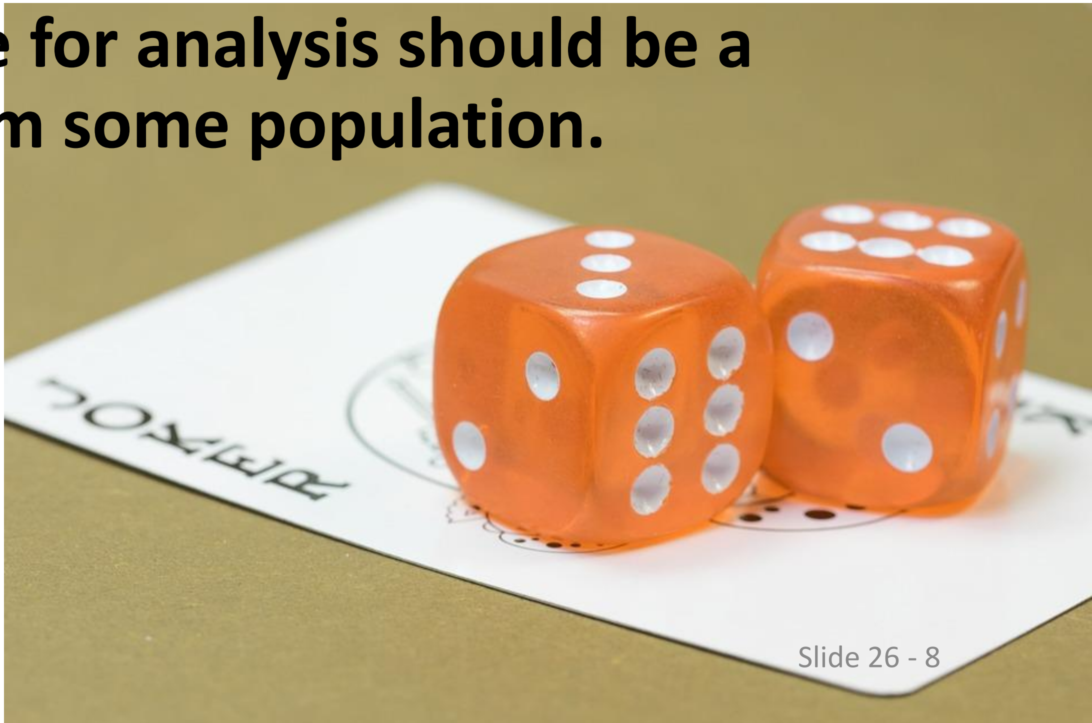

Assumptions and Conditions

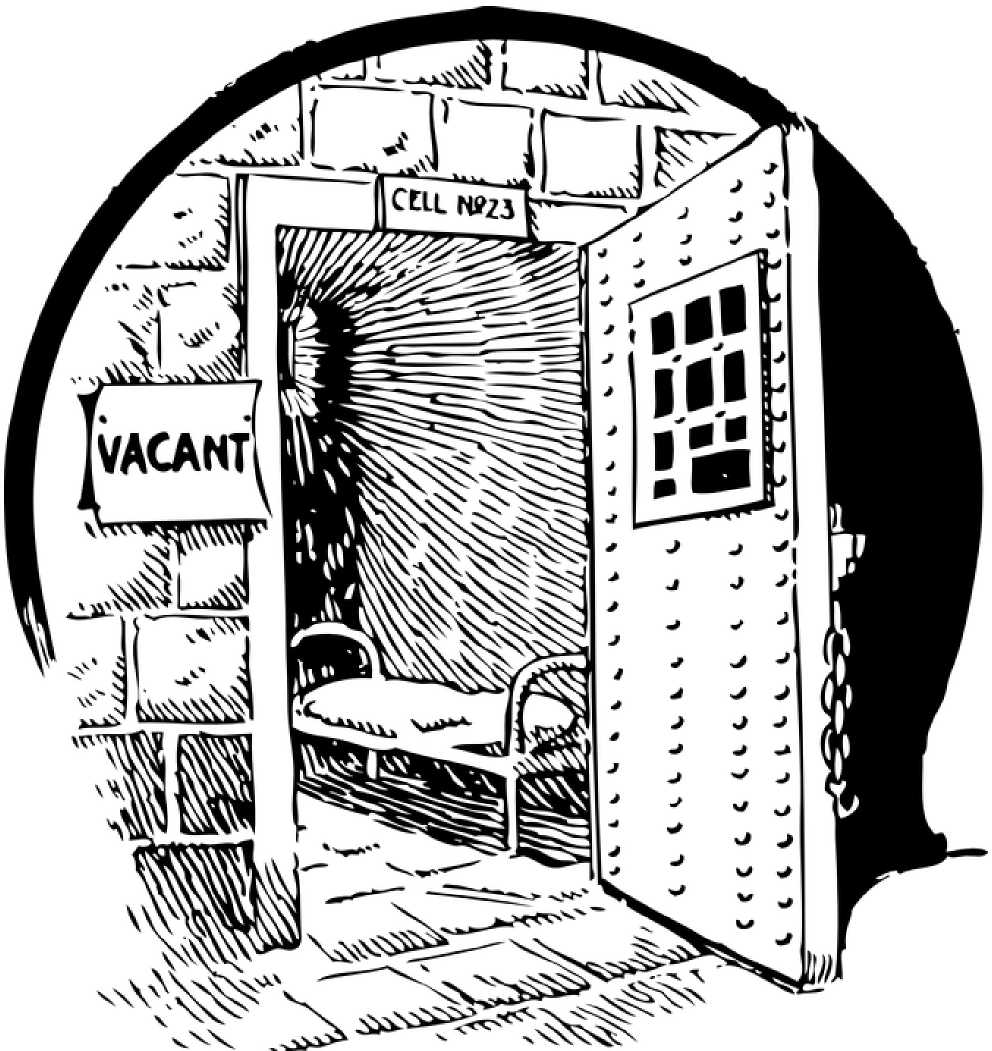

- ***Sample Size Assumption:*** We must have enough data for the methods to work.
  - ***Expected Cell Frequency Condition:*** We should expect to see at least 5 individuals in each cell.
    - This is similar to the condition that *np* and *nq* be at least 10 when we tested proportions.

Running the test

- How do I compare our count to the company?
  - Hypothesis
  - Conditions
  - Mechanics
  - Conclusion

Mechanics

$$\chi^2 = \sum_{all\ cells} \frac{(Obs - Exp)^2}{Exp}$$

- Since we want to examine how well the observed data reflect what would be expected, it is natural to look at the *differences* between the observed and expected counts (*Obs – Exp*).
- These differences are actually residuals, so we know that adding all of the differences will result in a sum of 0. That's not very helpful.
- We'll handle the residuals as we did in regression, by squaring them.
- To get an idea of the *relative* sizes of the differences, we will divide each squared difference by the expected count from that cell.

50% of 1094

Mechanics

$$\sum \frac{(Obs - Exp)^2}{Exp}$$

- Copy and complete the table

	Orange	Yellow	Brown	Total
Observed	503	293	298	1094
Expected	547	273.5	273.5	1094
O – E	-44	19.5	24.5	
$(O – E)^2$	1936	380.25	600.25	
$(O – E)^2/E$	3.539	1.39	2.194	7.123

X^2

χ^2 test statistic

Expected:

Equal counts (proportions) in each category

Calculations

df = categories − 1
3 − 1
df = 2

- The **chi-square models** are actually a family of distributions indexed by degrees of freedom (much like the *t*-distribution).
- The number of degrees of freedom for a goodness-of-fit test is $n - 1$, where n is the number of categories.

What to do with p-value

Due to a p-value of approx 0.025, which is lower than $\alpha = 0.05$, we reject the null. There is sufficient statistical evidence that the distribution of Reeces pieces is not as advertised.

p-value ≈ 0.025

Table entry for p is the point (χ^2) with probability p lying above it.

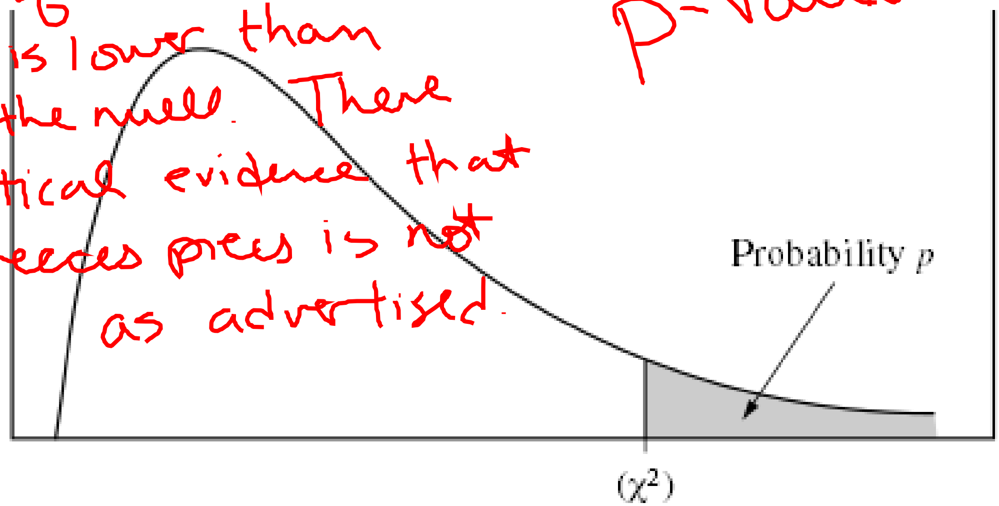

Table C χ^2 critical values

df	Tail probability p											
	.25	.20	.15	.10	.05	.025	.02	.01	.005	.0025	.001	.0005
1	1.32	1.64	2.07	2.71	3.84	5.02	5.41	6.63	7.88	9.14	10.83	12.12
2	2.77	3.22	3.79	4.61	5.99	7.38	7.82	9.21	10.60	11.98	13.82	15.20
3	4.11	4.64	5.32	6.25	7.81	9.35	9.84	11.34	12.84	14.32	16.27	17.73
4	5.39	5.99	6.74	7.78	9.49	11.14	11.67	13.28	14.86	16.42	18.47	20.00
5	6.63	7.29	8.12	9.24	11.07	12.83	13.39	15.09	16.75	18.39	20.51	22.11
6	7.84	8.56	9.45	10.64	12.59	14.45	15.03	16.81	18.55	20.25	22.46	24.10
7	9.04	9.80	10.75	12.02	14.07	16.01	16.62	18.48	20.28	22.04	24.32	26.02
8	10.22	11.03	12.03	13.36	15.51	17.53	18.17	20.09	21.95	23.77	26.12	27.87

Using your calculator...

One-Sided or Two-Sided?

- **The chi-square statistic is used only for testing hypotheses, not for constructing confidence intervals.**
- **If the observed counts don't match the expected, the statistic will be large—it can't be "too small."**
- **So the chi-square test is always one-sided.**
  - **If the calculated statistic value is large enough, we'll reject the null hypothesis.**

One-Sided or Two-Sided?

- **The mechanics may work like a one-sided test, but the interpretation of a chi-square test is in some ways *many*-sided.**
- **There are many ways the null hypothesis could be wrong.**
- **There's no direction to the rejection of the null model—all we know is that it doesn't fit.**

The Chi-Square Calculation

1. **Find the observed values**
2. **Find the expected values:**
   - Every model gives a hypothesized proportion for each cell.
   - The expected value is the product of the total number of observations times this proportion.
3. **Compute the residuals:** Once you have expected values for each cell, find the residuals, *Observed – Expected*.
4. **Square the residuals.**

The Chi-Square Calculation

5. **Compute the components.** Now find the components

$$\frac{(Observed - Expected)^2}{Expected}$$

for each cell.

6. **Find the sum of the components** (that's the chi-square statistic).
7. **Find the degrees of freedom.** It's equal to the number of cells minus one.
8. **Test the hypothesis.**
   - Use your chi-square statistic to find the P-value. (Remember, you'll always have a one-sided test.)
   - Large chi-square values mean lots of deviation from the hypothesized model, so they give small P-values.

Try One

- For 1000 shoppers donating blood at a mall, the frequencies of blood types were as shown in the table below. We want to know if the distribution of blood types is the same for customers in the mall as the general population. Consider this an SRS of all mall shoppers.

O	A	B	AB	TOTAL
465	394	96	45	1000

- In the general population, the blood type distribution is as follows:
  Type O = 45%, Type A = 40%, Type B = 10%, Type AB = 5%.

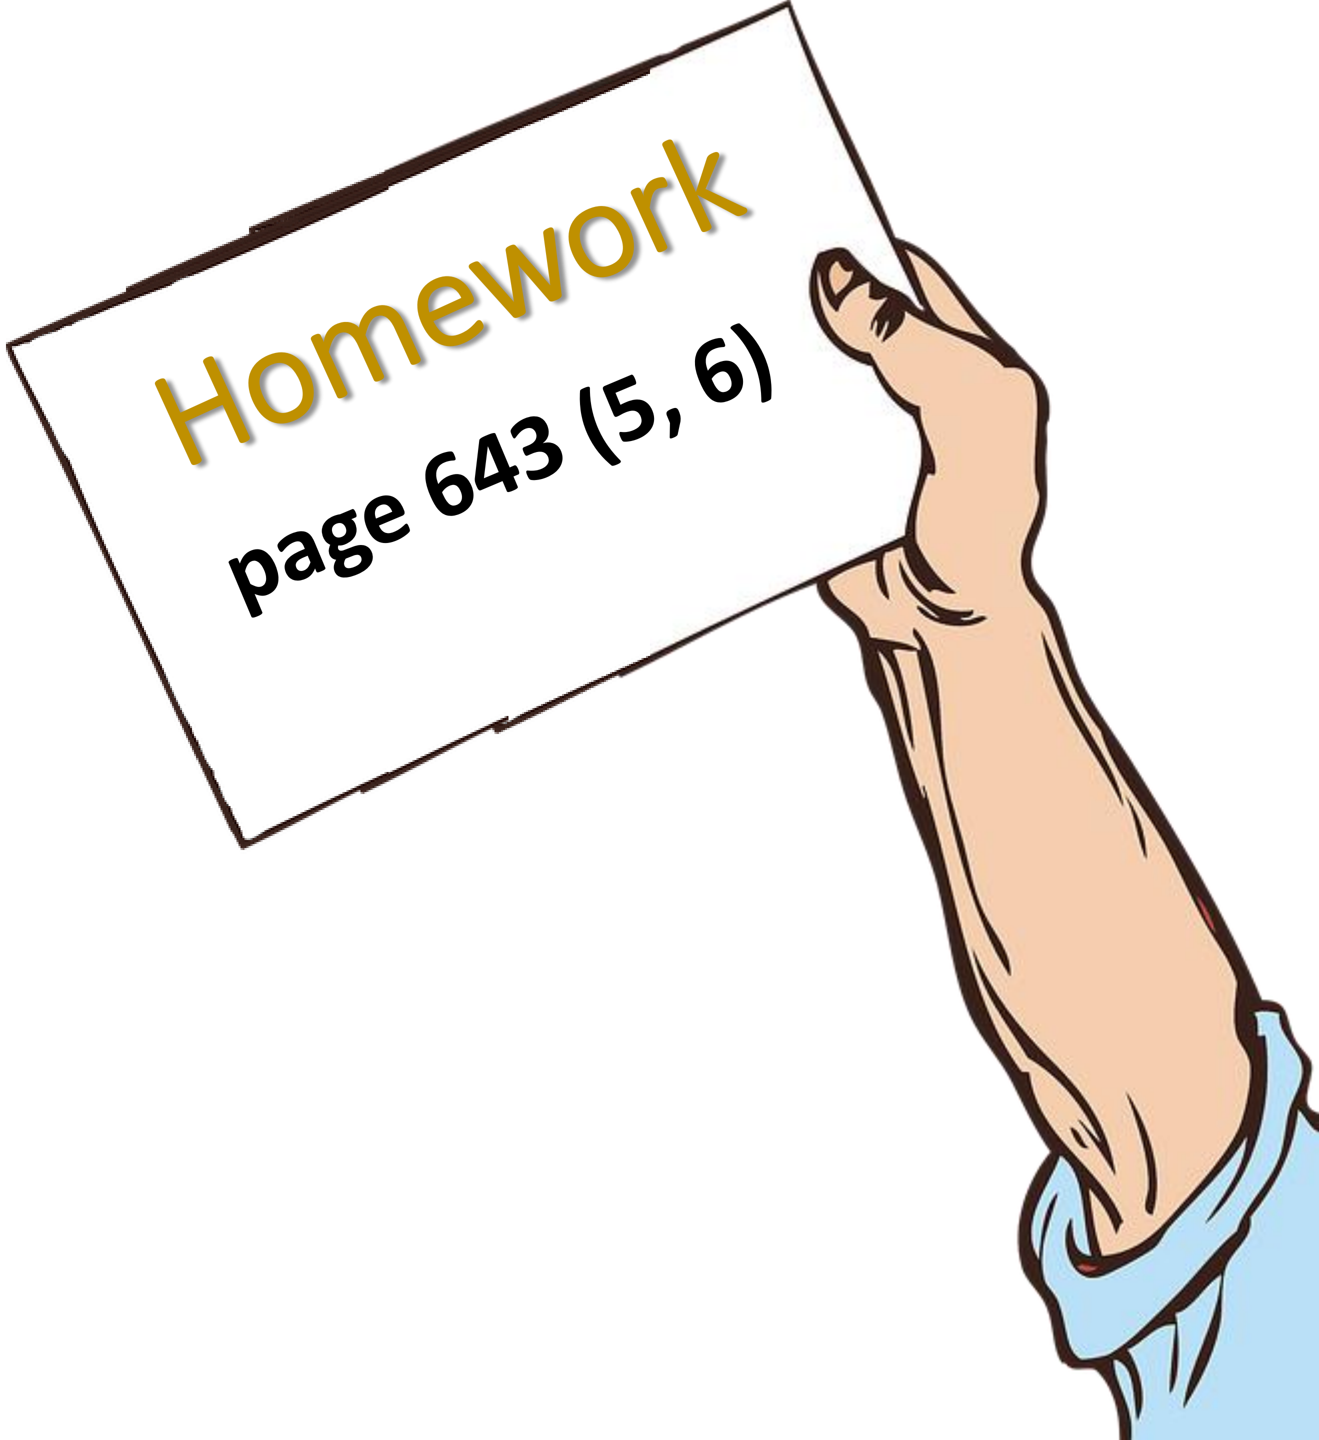
Homework
page 643 (5, 6)