EDP308: STATISTICAL LITERACY

The University of Texas at Austin, Fall 2020

RAZ: Rebecca A. Zárate, MA

Overview

- Goodness of Fit Test
  - Equal Proportions
  - Unequal Proportions
- Goodness of Fit Test Examples
  - Movie Genre
  - UT vs. USA Race-Ethnic Demographics
  - Skittles
- Police Killings in 2015
- Chi-Squared Goodness of Fit Test in R

Chi-Squared Goodness of Fit

How good is the fit?

- Let’s say I am trying to test some hypothesis about the US population.
  - I want to be sure to get an accurate representation of race and ethnicity.
- I take a sample of 100 UT students…

Will this sample be representative?
Will it have a good “fit”?

Goodness of Fit Test

- Chi-Squared Goodness of Fit Test is like the Chi-Squared Test of Independence, but now we are only have ONE (rather than two) categorical variables, and we want to know is if the proportion in each group is either:
  - Equal Proportions
    - Sex in the Population
      - Ex. 50% male, 50% female
  - Unequal Proportions
    - Age of the Work Force in USA
      - Ex. 36% 18-44 years, 26% 45-64 years, 13% 65+

Goodness of Fit Examples

- Do the proportions of admission applications to UT from different parts of the state match the proportions of people that live in that area?
- At a wedding, based on people's preferences, what proportions of songs should be waltzes, dance songs, and cumbias?
- Are there equal proportions of Skittle colors in a bag?
- Are there equal proportions of Men and Women at UT?

Goodness of Fit Hypotheses

- As usual, our null hypothesis is that the proportions are all equal (or equal to some known proportions)

$$H_0\colon p_1 = p_2 = p_3$$

$$H_0\colon The\ proportions\ are\ the\ same$$
$$(or\ equal\ to\ a\ known\ pattern)\, for\ each\ level/group\ of\ the\ variable.$$

- The alternative hypothesis is that the proportions are not the same (or do not follow the known pattern) for each level/group of the variable:

$$H_1\colon p_1 \neq p_2 \neq p_3$$

$$H_1\colon The\ proportions\ are\ not\ the\ same\ for\ each\ level$$
$$of\ the\ variable\ (or\ do\ not\ match\ a\ known\ pattern).$$

Try it. Movie Genre

- You go out and ask 120 Netflix and Chill people to see if there is a preference for genre of movie. Use the goodness of fit test at $\alpha = .05$ to test this.

What would the EXPECTED values be for each genre?

Favorite Genre	Observed
Action	32
Comedy	24
Romance	35
Horror	29
Total	120

$$H_0\text{: } p_{Action} = p_{Comedy} = p_{Romance} = p_{Horror} = .25$$
$$H_1\text{: } The\ proportions\ are\ not\ the\ same\ for\ each\ movie\ genre.$$

Try it. Movie Genre

- If we assuming the null is true (genre preferences are all equal), then we would expect equal frequency for all the genres.
  - 120 (responses)/4 (genres) = 30

Favorite Genre	Observed	Expected
Action	32	30
Comedy	24	30
Romance	35	30
Horror	29	30
Total	120	120

Try it. Movie Genre

Step 1:

$$H_0: p_{Action} = p_{Comedy} = p_{Romance} = p_{Horror} = .25$$
$$H_1: The\ proportions\ are\ not\ the\ same\ for\ each\ genre$$

Step 2:

$$\alpha = .05$$

Step 3:

$$df = 4 - 1 = 3$$
$$df = 3$$

Step 4:

$$\chi^2_{crit} = 7.81$$

Step 5,6: Compute Test Statistic and Conclusions

Favorite Genre	Observed	Expected
Action	32	30
Comedy	24	30
Romance	35	30
Horror	29	30
Total	120	120

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

$$\chi^2 = \frac{(32-30)^2}{30} + \frac{(24-30)^2}{30} + \frac{(35-30)^2}{30} + \frac{(29-30)^2}{30} \approx \mathbf{2.2}$$

Using $\alpha = .05$ and $df = 3$, our $\chi^2_{crit} = 7.81$. Because our χ^2_{stat} is not past our χ^2_{crit}, we fail to reject H_0.

People do tend to see movies in equal proportions.

Try it. UT and USA

Since we use college students for a lot of social science research, I want to know if the convenient UT sample I took is a "good fit" to represent the USA. Conduct a Chi-Squared Goodness of Fit test on the data below using $\alpha = .05$.

Race/Ethnicity	Observed	Expected
White	41	61
Black	4	13
Latinx	21	16
Asian	19	5
Other	15	5

Try it. UT and USA

Step 1:

$$H_0\text{: } A \text{ sample of UT students matches the national averages of race and ethnicity}$$
$$H_1\text{: } A \text{ sample of UT students does not match the national averages of race and ethnicity}$$

Step 2:

$$\alpha = .05$$

Step 3:

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

Step 4:

$$df = 4 \text{ and } \alpha = .05,\ \chi^2_{crit} = 9.49$$

Step 5,6: Compute Test Statistic and Conclusions

Race/Ethnicity	Observed	Expected
White	41	61
Black	4	13
Latinx	21	16
Asian	19	5
Other	15	5

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

$$\chi^2 = \frac{(41-61)^2}{61} + \frac{(4-13)^2}{13} + \frac{(21-16)^2}{16} + \frac{(19-5)^2}{5} + \frac{(15-5)^2}{5} \approx \mathbf{73.55}$$

Using $\alpha = .05$ and $df = 4$, our $\chi^2_{stat} = 73.55$. Because our χ^2_{stat} is past our χ^2_{crit}, we reject H_0. The UT sample does not match the USA race/ethnic population.

Try it. Skittles

Test whether all the colors (flavors) of Skittles are present in equal proportions using $\alpha = .05$.

Color (Flavor)	In this bag	Expected
Red	80	73
Orange	87	73
Yellow	59	73
Green	70	73
Purple	60	73
Total	**365**	**365**

Step 1:

$$H_0\text{: The colors of Skittles are found in equal proportions.}$$
$$H_1\text{: The colors of Skittles are not found in equal proportions.}$$

Try it. Skittles

Step 2:

$$\alpha = .05$$

Step 3:

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

Step 4:

$$df = 5 - 1 = 4$$

With $df = 3$, and $\alpha = .05$, we find $\chi^2_{crit} = 9.49$

Try it. Skittles

Step 5:

Color (Flavor)	In this bag	Expected
Red	80	73
Orange	87	73
Yellow	59	73
Green	70	73
Purple	60	73
Total	**365**	**365**

$$\chi^2 = \frac{(80-73)^2}{73} + \frac{(87-73)^2}{73} + \frac{(59-73)^2}{73} + \frac{(70-73)^2}{73} + \frac{(60-73)^2}{73} \approx 8.47$$

Our $\chi^2_{stat} = 8.47$ and our $\chi^2_{crit} = 9.49$. Because our χ^2_{stat} does not pass χ^2_{crit}, we fail to reject the null hypothesis. There is no reason to suggest Skittles colors are not represented in equal proportions.

Police Killings in 2015

Police Killings in 2015

- For decades, we have heard of cases when a police office kills someone in the line of duty.
  - Sometimes these actions are justified, and sometimes they are not.
- Some hold the belief that certain minorities are targeted more frequently than the White majority.
  - Others dispute this and believe no racial or ethnic group is disproportionally targeted (NULL hypothesis).
- Let's use statistics to see if we can provide some evidence one way or the other.

Police Killings in 2015

- State the Hypotheses:
  - H_0: Police killings occur in equal proportions to the racial and ethnic demographics of the USA
    - Asian = 6%, Black = 13%, Latinx = 16%,

    Native American = 1%, (Non-Latinx) White = 64%
  - H_1: Police killings DO NOT occur in equal proportions to the racial and ethnic demographics of the USA
- $\alpha = .05$
- Using a Chi-Squared Goodness of Fit Test
- df = 5 (groups) – 1 = 4
- $\chi^2_{crit} = 9.49$

Police Killings in 2015 in R

Using data.

```
#########################################
############# Chi-Squared ###############
######## Goodness of Fit Test ###########
#########################################

# Reading in a data set of police killings from 2015
police_killings <- read.csv("police_killings_2015.csv")

# Filtering out cases where the race-ethnicity is Unknown or Native American
police_killings <- filter(police_killings, police_killings$raceethnicity != "Unknown")

police_killings_table <- table(police_killings$raceethnicity)

# Because we have an expected pattern for race-ethnicity in the USA, we need to tell R
# what those proportions are. Note: The proprotions must add up to 1. There is some rounding with these numbers
# so they have been slightly tweeked but not by more than 1% for any group
# The order: Asian = 6%, Black = 13%, Latinx = 16%, Native American = 1%, (Non-Latinx) White = 64%
expected_proportions <- c(.06, .13, .16, .01, .64)

# We then put the UT data and the expected proportions into the "chisq.test()" function
police_killings_chi_squared <- chisq.test(police_killings_table, p = expected_proportions)
```

Data Sources: https://github.com/fivethirtyeight/data/tree/master/police-killings
https://www.theguardian.com/us-news/ng-interactive/2015/jun/01/the-counted-police-killings-us-database
https://fivethirtyeight.com/features/where-police-have-killed-americans-in-2015/

Police Killings in 2015 in R

- We see that the chi-squared test is significant. We can reject the null hypothesis. The number of police killings for a certain race-ethnicity, do not match the proportions of those race-ethnicities in the USA.

```
> police_killings_chi_squared

	Chi-squared test for given probabilities

data:  police_killings_table
X-squared = 119.99, df = 4, p-value < 2.2e-16
```

- We can also look at the information stored in the "police_killings_chi_squared" object by clicking on it.

Police Killings in 2015 in R

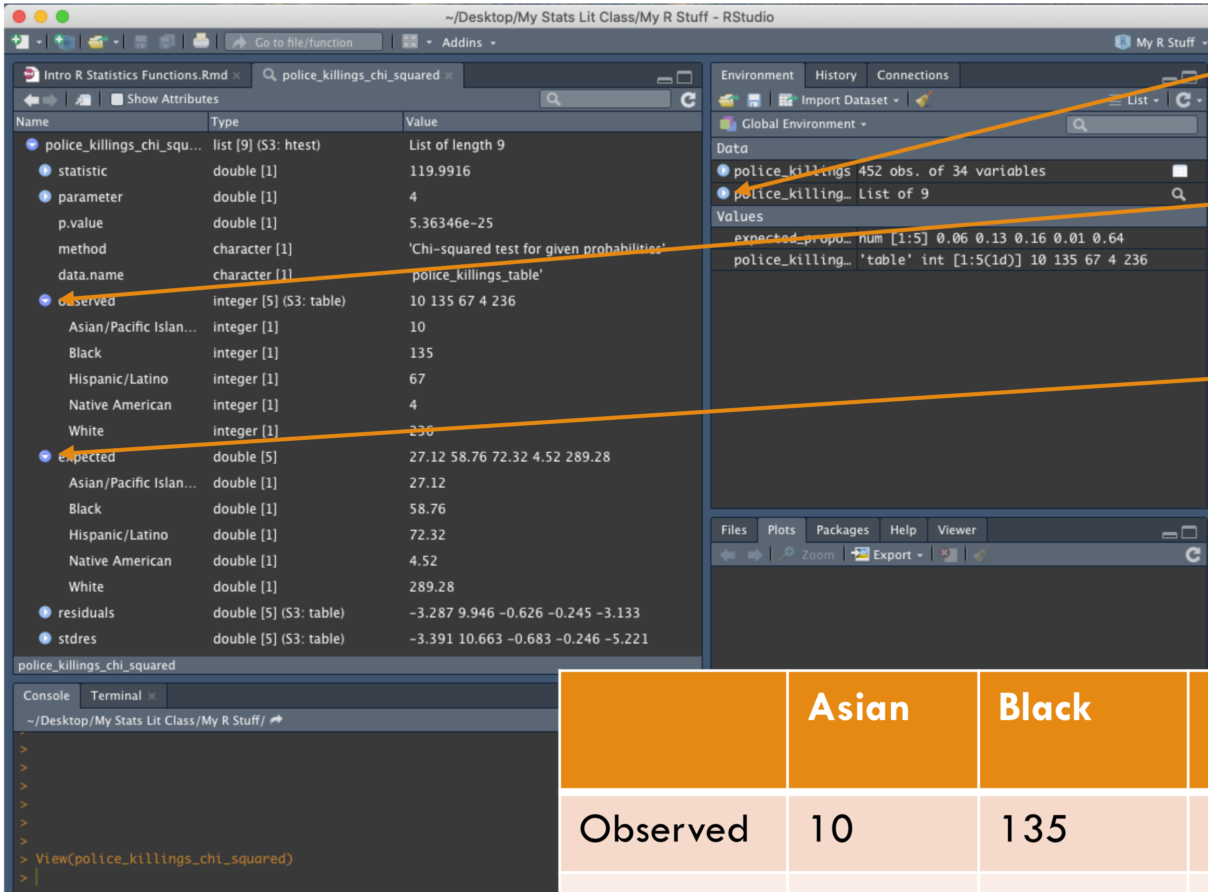

Click here to see the "police_killings_chi_squared" object.

Click here to see the observed values, i.e. the number of people killed by police for each race-ethnicity

Click here to see the expected values, i.e. the number of people we would expect to be killed by police for each race-ethnicity if the null were true.

	Asian	Black	Latinx	Native American	White
Observed	10	135	67	4	236
Expected	27	59	72	5	289
Difference	-17	+76	-5	-1	-53

Chi-Squared Goodness of Fit Test in R

Chi-Squared Goodness of Fit Test in R

```
########################################
############ Chi-Squared ###############
######## Goodness of Fit Test ###########
########################################

# Movie Genres
movie_genre <- as.table(rbind(c(32, 24, 35, 29)))
dimnames(movie_genre) <- list(Status = c("Action", "Comedy", "Romance", "Horror"))

chisq.test(movie_genre)

# UT vs USA Racial-Ethnic Demographics (UNEQUAL proportions)
ut <- as.table(rbind(c(41, 4, 21, 19, 15)))
dimnames(ut) <- list(Status = c("White", "Black", "Latinx", "Asian", "Other"))

# Because we have an expected pattern for race-ethnicity in the USA, we need to tell R
# what those proportions are
expected_proportions <- c(.61, .13, .16, .05, .05)

# We then put the UT data and the expected proportions into the "chisq.test()" function
chisq.test(ut, p = expected_proportions)

# Skittles
skittles <- as.table(rbind(c(80, 87, 59, 70, 60)))
dimnames(skittles) <- list(Colors = c("Red", "Orange", "Yellow", "Green", "Purple"))

chisq.test(skittles)
```

Using summary data.

Chi-Squared Goodness of Fit Test in R Output

- Movies genres are watched equally.
- Our UT sample is does not “fit” the USA population.
- Skittles colors are equally represented in a bag.

```
> chisq.test(movie_genre)

	Chi-squared test for given probabilities

data:  movie_genre
X-squared = 2.2, df = 3, p-value = 0.5319

> # We then put the UT data and the expected proportions into the "chisq.test()" function
> chisq.test(ut, p = expected_proportions)

	Chi-squared test for given probabilities

data:  ut
X-squared = 73.551, df = 4, p-value = 4.035e-15

> chisq.test(skittles)

	Chi-squared test for given probabilities

data:  skittles
X-squared = 8.4663, df = 4, p-value = 0.07592
```

Using summary data.