Chi-Square Tests

Nathaniel E. Helwig

Associate Professor of Psychology and Statistics University of Minnesota



October 17, 2020

Copyright \bigodot 2020 by Nathaniel E. Helwig

Nathaniel E. Helwig (Minnesota)

Chi-Square Tests

1. Goodness of Fit

2. Tests of Association (for 2-way Tables)

3. Conditional Association Tests (for 3-way Tables)

Table of Contents

1. Goodness of Fit

2. Tests of Association (for 2-way Tables)

3. Conditional Association Tests (for 3-way Tables)

A Primer on Categorical Data Analysis

In the previous chapter, we looked at inferential methods for a single proportion or for the difference between two proportions.

In this chapter, we will extend these ideas to look more generally at contingency table analysis.

All of these methods are a form of "categorical data analysis", which involves statistical inference for nominal (or categorial) variables.

Categorical Data with J > 2 Levels

Suppose that X is a categorical (i.e., nominal) variable that has J possible realizations: $X \in \{0, ..., J-1\}$. Furthermore, suppose that

$$P(X=j) = \pi_j$$

where π_j is the probability that X is equal to j for $j = 0, \ldots, J - 1$.

Assume that the probabilities satisfy $\sum_{j=0}^{J-1} \pi_j = 1$, so that $\{\pi_j\}_{j=0}^{J-1}$ defines a valid probability mass function for the random variable X.

- If J = 2, this is a Bernoulli distribution
- If J > 2, this is a multinomial distribution

Testing Hypotheses about Multiple Probabilities

Suppose we have a sample of data x_1, \ldots, x_n where $x_i \stackrel{\text{iid}}{\sim} F$ with F denoting the probability mass function defined by $\{\pi_j\}_{j=0}^{J-1}$.

We want to test a null hypothesis of the form $H_0: \pi_j = \pi_{j0} \ \forall j$ versus the alternative hypothesis $H_1: (\exists j)(\pi_j \neq \pi_{j0})$.

- The symbol \forall should be read as "for all"
- The symbol \exists should be read as "there exists"

 H_0 states that the probability for the *j*-th category is equal to π_{j0} for all $j = 0, \ldots, J - 1$, and H_1 states that there exists at least one category where the probability for the *j*-th category is not equal to π_{j0} .

Observed and Expected Frequencies

Given an iid sample of n observations of the random variable X, the observed frequency for the j-th category is given by

$$f_j = \sum_{i=1}^n I(x_i = j)$$

for j = 0, ..., J - 1, which is the number of x_i belonging to category j.

Given an iid sample of n observations of the random variable X, the expected frequency for the j-th category is given by

$$m_j = n\pi_j$$

for $j = 0, \ldots, J - 1$, which is the sample size n multiplied by π_j .

Test Statistic

To test $H_0: \pi_j = \pi_{j0} \ \forall j$ versus $H_1: (\exists j)(\pi_j \neq \pi_{j0})$, define

$$X^{2} = \sum_{j=0}^{J-1} \frac{(f_{j} - m_{j0})^{2}}{m_{j0}}$$

where $m_{j0} = n\pi_{j0}$ is the expected frequency assuming that H_0 is true.

If the observed frequencies are far from what would be expected assuming H_0 is true, then the value of X^2 would be relatively large.

As the sample size $n \to \infty$, the X^2 test statistic approaches a χ^2 distribution with J-1 degrees of freedom (Pearson, 1900).

Example 1: Vaping Prevalence

Suppose a researcher is interested in studying the prevalence of vaping among college students in the United States.

The researcher asked a random sample of n = 1000 college students the typical number of pods they vape per day, and finds the following:

# Pods	Observed Frequency	Expected Frequency
0	780	830
0-1	140	110
1-2	60	50
> 2	20	10

Expected frequencies are based on last year's data, which found: $\pi_{00} = 0.83$, $\pi_{10} = 0.11$, $\pi_{20} = 0.05$, and $\pi_{30} = 0.01$.

Example 1: Vaping Prevalence (continued)

To test the null hypothesis $H_0: \pi_j = \pi_{j0} \forall j$ versus the alternative hypothesis $H_1: (\exists j)(\pi_j \neq \pi_{j0})$, the X^2 test statistic is

$$X^{2} = \frac{(780 - 830)^{2}}{830} + \frac{(140 - 110)^{2}}{110} + \frac{(60 - 50)^{2}}{50} + \frac{(20 - 10)^{2}}{10} = 23.19387$$

Comparing this to a χ^2 distribution with three degrees of freedom

$$p = P(\chi_3^2 > 23.19387) = 0.00003679 = 3.679 \times 10^{-5}$$

Thus, we have reason to suspect that the prevalence of vaping this year is significantly different from last year.

Nathaniel E. Helwig (Minnesota)

Example 1: Vaping Prevalence (in R)

Chi-squared test for given probabilities

data: f
X-squared = 23.194, df = 3, p-value = 3.679e-05

Table of Contents

1. Goodness of Fit

2. Tests of Association (for 2-way Tables)

3. Conditional Association Tests (for 3-way Tables)

Joint Distribution for Two Categorical Variables

Suppose that A and B are both categorical (i.e., nominal) variables with $A \in \{1, ..., J\}$ and $B \in \{1, ..., K\}$.

Furthermore, suppose that...

•
$$P(A = j) = \pi_j$$
. for $j = 1, ..., J$ (with $\sum_{j=1}^J \pi_j = 1$)
• $P(B = k) = \pi_{k}$ for $k = 1, ..., K$ (with $\sum_{k=1}^K \pi_k = 1$)

And assume that the probability of observing the joint event is given by

$$P(A=j\cap B=k)=\pi_{jk}$$

where the joint probabilities satisfy $\sum_{j=1}^{J} \sum_{k=1}^{K} \pi_{jk} = 1$.

2×2 Contingency Table

Given an iid sample of n observations $(a_i, b_i) \stackrel{\text{iid}}{\sim} F$ from the joint probability distribution F, the $J \times K$ table that contains the observed frequency for each combination of A and B is referred to as a contingency table, which is also known as a <u>cross tabulation</u>.

	B = 1	B=2		B = k	•••	B = K	Row Totals
A = 1	f_{11}	f_{12}	• • •	f_{1k}	•••	f_{1K}	f_1 .
A=2	f_{21}	f_{22}	•••	f_{2k}	•••	f_{2K}	f_2 .
:	÷	÷	·	÷	·	÷	÷
A = j	f_{j1}	f_{j2}	•••	f_{jk}	•••	f_{jK}	$f_j.$
:	:	:	۰.	÷	·	÷	÷
A = J	f_{J1}	f_{J2}		f_{Jk}	•••	f_{JK}	f_J .
Column Totals	$f_{\cdot 1}$	$f_{\cdot 2}$	•••	$f_{\cdot k}$	•••	$f_{\cdot K}$	$f_{\cdot\cdot} = n$

2×2 Contingency Table (continued)

In the previous table, f_{jk} denotes the number of observations that are cross-classified in category j of variable A and category k of variable B

$$f_{jk} = \sum_{i=1}^{n} I(a_i = j) I(b_i = k)$$

for all $j, k \in \{1, ..., J\} \times \{1, ..., K\}.$

The row and column totals are the marginal observed frequencies for variables A and B, which are defined as

$$f_{j.} = \sum_{i=1}^{n} I(a_i = j) = \sum_{k=1}^{K} f_{jk}$$
$$f_{.k} = \sum_{i=1}^{n} I(b_i = k) = \sum_{j=1}^{J} f_{jk}$$

for all $j \in \{1, \ldots, J\}$ and all $k \in \{1, \ldots, K\}$.

Test of Statistical Independence

To test if A and B are independent of one another, we want to test $H_0: \pi_{jk} = \pi_j . \pi_{\cdot k} \ \forall j, k$ versus $H_1: (\exists j, k) (\pi_{jk} \neq \pi_j . \pi_{\cdot k}).$

Assuming that H_0 is true, the expected number of observations cross-classified in cell (j, k) of the contingency table would be

 $m_{jk} = n\pi_j.\pi_{k}$

which is the sample size n multiplied by the marginal probabilities.

Use $\hat{m}_{ik} = n\hat{\pi}_i \cdot \hat{\pi}_{\cdot k}$ in practice (because π_i and $\pi_{\cdot k}$ are unknown)

Test Statistic

Given the estimates of the expected cell counts (assuming H_0 is true), we can use the chi-square test statistic

$$X^{2} = \sum_{j=1}^{J} \sum_{k=1}^{K} \frac{(f_{jk} - \hat{m}_{jk})^{2}}{\hat{m}_{jk}}$$

to test the null hypothesis of independence between A and B.

Assuming that H_0 is true, we have that $X^2 \sim \chi^2_{(J-1)(K-1)}$ as $n \to \infty$, which is Pearson's chi-square test for association (Pearson, 1900)

The chi-square approximation is because, assuming H_0 is true, we have

$$z_{jk} = \frac{f_{jk} - \hat{m}_{jk}}{\sqrt{\hat{m}_{jk}}} \stackrel{d}{\to} N(0, 1)$$

as the sample size $n \to \infty$.

Example 2: Race and The Death Penalty

This contingency table is from Table 4 of Radelet and Pierce (1991), which cross-classifies individuals by race and death penalty sentence:

Death Penalty						
Defendant	Yes	No	Total			
White	53	430	483			
Black	15	176	191			
Total	68	606	674			

Suppose that we want to test the null hypothesis that race and death penalty sentence are independent.

Example 2: Race and The Death Penalty (continued)

The marginal probability estimates for the rows are

 $\hat{\pi}_{1.} = 483/674 = 0.7166172$ and $\hat{\pi}_{2.} = 191/674 = 0.2833828$

and the marginal probability estimates for the columns are

 $\hat{\pi}_{.1} = 68/674 = 0.1008902$ and $\hat{\pi}_{.2} = 606/674 = 0.8991098$

This implies that the null hypothesized (estimates of the) expected frequency for each cell is

	Death Penalty				
Defendant	Yes	No			
White	48.72997	434.27			
Black	19.27003	171.73			

Example 2: Race and The Death Penalty (continued)

The Pearson's chi-square test statistic is given by

$$X^{2} = \frac{(53 - 48.72997)^{2}}{48.72997} + \frac{(430 - 434.27)^{2}}{434.27} + \frac{(15 - 19.27003)^{2}}{19.27003} + \frac{(176 - 171.73)^{2}}{171.73}$$

= 1.468519

Comparing this to a χ_1^2 distribution, the p-value for testing H_0 is

$$p = P(\chi_1^2 > 1.468519) = 0.2255796$$

Don't have sufficient evidence to reject the null hypothesis that race and death penalty sentence are independent.

Nathaniel E. Helwig (Minnesota)

Example 2: Race and The Death Penalty (in R)

We can confirm this result using the chisq.test function in R:

```
> xtab <- matrix(c(53, 15, 430, 176), 2, 2)
> colnames(xtab) <- c("Yes", "No")
> rownames(xtab) <- c("White", "Black")
> xtab
           Yes No
White 53 430
Black 15 176
> chisq.test(xtab, correct = FALSE)
Pearson's Chi-squared test
```

```
data: xtab
X-squared = 1.4685, df = 1, p-value = 0.2256
```

Relation to Testing Proportion Difference

The result of Pearson's chi-squared test is the exact same as the asymptotic test for the difference in proportions that was conducted in the previous chapter (via the prop.test function).

For a 2×2 contingency table, it can be shown that the X^2 test statistic is identical to the Z^2 test statistic that was used for the asymptotic test of the difference between two proportions.

Testing the null hypothesis $H_0: \pi_{jk} = \pi_j . \pi_{\cdot k} \; \forall j, k$ is equivalent to testing the null hypothesis $H_0: \pi_{1.} = \pi_2$.

• A and B being independent implies equal probabilities of success

Table of Contents

1. Goodness of Fit

2. Tests of Association (for 2-way Tables)

3. Conditional Association Tests (for 3-way Tables)

Three-Way Contingency Tables

A three-way contingency table cross-classifies observations on three different categorical variables, such as the below table.

Previously, we looked at a 2-way table (Defendant's Race by Death Penalty), however these data could actually be arranged into a 3-way table where the third variable is the Victim's Race.

Note that all of these defendants were on trial for committing multiple homicides, so we can look at how both the defendant's and victim's race affects the probability of receiving a death penalty sentence.

Example of Three-Way Table

		Death Penalty			
Victim	Defendant	Yes	No	Total	
White	White	53	414	467	
winte	Black	11	37	48	
Block	White	0	16	16	
DIACK	Black	4	139	143	
Tetel	White	53	430	483	
Total	Black	15	176	191	

A = Defendant's Race B = Death Penalty Verdict C = Victim's Race

Marginal and Partial Tables

The previous two-way table that we looked at was a "marginal table", which is formed by aggregating the three-way table across the levels of one variable (in this case Victim's Race).

We could also look at the two relevant "partial tables", which are two-way slices of the three-way table.

White Victim	Death	Penalty		Black Victim	Death	Penalty	
Defendant	Yes	No	Total	Defendant	Yes	No	Total
White	53	414	467	White	0	16	16
Black	11	37	48	Black	4	139	143
Total	64	451	515	Total	4	155	159

Marginal and Conditional Independence

Variables A and B are said to be <u>marginally independent</u> if they are independent after aggregating across the levels of a third variable C.

Variables A and B are said to be conditionally independent given C if they are independent at all levels of the third variable C.

Marginal independence does not imply conditional independence. Conditional independence does not imply marginal independence.

• Simpson's paradox: when marginal and conditional effects differ

Example 3: Conditional Independence Test

Previously, we tested the null hypothesis that A (Defendant's Race) and B (Death Sentence Verdict) are marginally independent after collapsing across both of the levels of C (Victim's Race).

• No evidence to reject the null hypothesis of independence

We will test the conditional independence between A and B given C.

- $H_0: \pi_{jk(\ell)} = \pi_{j \cdot (\ell)} \pi_{\cdot k(\ell)} \ \forall j, k, \ell$
- $H_1: (\exists j, k, \ell)(\pi_{jk(\ell)} \neq \pi_{j \cdot (\ell)} \pi_{\cdot k(\ell)})$
- $\ell \in \{1,2\}$ denotes the Victim's Race

Example 3: White Victim

When the victim is white, the estimated probabilities of receiving the death penalty are

 $\hat{\pi}_{1\cdot(1)} = 53/467 = 0.1134904$ and $\hat{\pi}_{2\cdot(1)} = 11/48 = 0.2291667$

Note that the probability of receiving the death penalty is twice as high for black defendants when the victim is white.

```
> white.victim <- matrix(c(53, 11, 414, 37), 2, 2)
> colnames(white.victim) <- c("Yes", "No")
> rownames(white.victim) <- c("White", "Black")
> chisq.test(white.victim, correct = FALSE)
Pearson's Chi-squared test
data: white.victim
```

```
X-squared = 5.3518, df = 1, p-value = 0.0207
```

Example 3: Black Victim

When the victim is black, the estimated probabilities of receiving the death penalty are

 $\hat{\pi}_{1\cdot(2)} = 0/16 = 0$ and $\hat{\pi}_{2\cdot(2)} = 4/143 = 0.02797203$

```
> black.victim <- matrix(c(0, 4, 16, 139), 2, 2)
> colnames(black.victim) <- c("Yes", "No")
> rownames(black.victim) <- c("White", "Black")
> chisq.test(black.victim, correct = FALSE)
Pearson's Chi-squared test
```

```
data: black.victim
X-squared = 0.4591, df = 1, p-value = 0.498
```

```
Warning message:
In chisq.test(black.victim, correct = FALSE) :
   Chi-squared approximation may be incorrect
```

Nathaniel E. Helwig (Minnesota)

Example 3: Conclusions

Regardless of whether we can trust the results for the black victims...

It seems apparent that the Defendant's Race and the Death Penalty Sentence are **not** conditionally independent given the Victim's Race.

Evidence suggests that black defendants are more likely than white defendants to receive the death penalty if the homicide victim is white.

References

- Pearson, K. (1900). On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *Philosophical Magazine* 50(302), 157–175.
- Radelet, M. L. and G. L. Pierce (1991). Choosing those who will die: Race and the death penalty in Florida. *Florida Law Review* 43(1), 1–34.