Lesson 40: How well do the observed data correspond to an idealized situation (theory)?

Does your zodiac predict how successful you will be later in life?

Scenario: Fortune magazine collected the zodiac signs of 256 heads of the largest 400 companies.

Births	Sign
23	Aries
20	Taurus
18	Gemini
23	Cancer
20	Leo
19	Virgo
18	Libra
21	Scorpio
19	Sagittarius
22	Capricorn
24	Aquarius
29	Pisces

What do you observe in the table?

Births	Sign
23	Aries
20	Taurus
18	Gemini
23	Cancer
20	Leo
19	Virgo
18	Libra
21	Scorpio
19	Sagittarius
22	Capricorn
24	Aquarius
29	Pisces

If births were distributed uniformly across the year, we would expect about 1/12 of them to occur under each sign.

256/12 = 21.3 births per sign

How closely do the observed numbers of births per sign fit this ideal model?

A hypothesis test used to address this question is called a test of "goodness-of-fit."

Goodness-of-fit Test

not %, proportions, measurements

Conditions:

1. **Counted data condition**: data must be counts for the categories.
2. **Independence assumption**: counts in the cells should be independent of each other.
3. **Randomization**: counted individuals should be a random sample
4. **Sample Size** (Expected cell frequency condition): at least five counts in each cell.

Degrees of freedom = Categories -1

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

H_0: the model fits

H_a: the model doesn't fit (some other model fits)

Chi-squared models

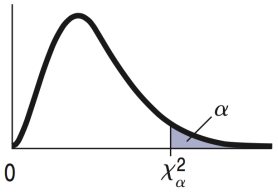

- Like student's t models
- different model for different degrees of freedom
- chi-squared tests are always one-tailed
- If χ^2 is large enough, we'll reject the null hypothesis

We have counts of 256 executives in 12 zodiac sign categories. The natural null hypothesis is that birth dates of executives are divided equally among all the zodiac signs. The test statistic looks at how closely the observed data match this idealized situation.

Question: Are zodiac signs of CEOs distributed uniformly?

I want to know whether births of successful people are uniformly distributed across the signs of the zodiac. I have counts of 256 Fortune 400 executives, categorized by their birth sign.

H_0: Births are uniformly distributed over zodiac signs.[2]

H_A: Births are not uniformly distributed over zodiac signs.

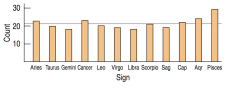

The bar chart shows some variation from sign to sign, and Pisces is the most frequent. But it is hard to tell whether the variation is more than I'd expect from random variation.

- ✓ **Counted Data Condition:** I have counts of the number of executives in 12 categories.
- ✓ **Independence Assumption:** The birth dates of executives should be independent of each other.
- ✓ **Randomization Condition:** This is a convenience sample of executives, but there's no reason to suspect bias.
- ✓ **Expected Cell Frequency Condition:** The null hypothesis expects that 1/12 of the 256 births, or 21.333, should occur in each sign. These expected values are all at least 5, so the condition is satisfied.

The conditions are satisfied, so I'll use a χ^2 model with $12 - 1 = 11$ degrees of freedom and do a **chi-square goodness-of-fit test.**

The expected value for each zodiac sign is 21.333.

$$\chi^2 = \sum \frac{(Obs - Exp)^2}{Exp} = \frac{(23 - 21.333)^2}{21.333} + \frac{(20 - 21.333)^2}{21.333} + \cdots$$

$= 5.094$ for all 12 signs.

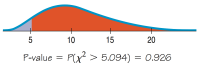

P-value = $P(\chi^2 > 5.094) = 0.926$

The P-value of 0.926 says that if the zodiac signs of executives were in fact distributed uniformly, an observed chi-square value of 5.09 or higher would occur about 93% of the time. This certainly isn't unusual, so I fail to reject the null hypothesis, and conclude that these data show virtually no evidence of nonuniform distribution of zodiac signs among executives.

How big is big? When we calculated χ^2 for the zodiac sign example, we got 5.094. That value would have been big for z or t, leading us to reject the null hypothesis. Not here. Were you surprised that $\chi^2 = 5.094$ had a huge P-value of 0.926? What *is* big for a χ^2 statistic, anyway?

Think about how χ^2 is calculated. In every cell, any deviation from the expected count contributes to the sum. Large deviations generally contribute more, but if there are a lot of cells, even small deviations can add up, making the χ^2 value larger. So the more cells there are, the higher the value of χ^2 has to get before it becomes noteworthy. For χ^2, then, the decision about how big is big depends on the number of degrees of freedom.

Unlike the Normal and t families, χ^2 models are skewed. Curves in the χ^2 family change both shape and center as the number of degrees of freedom grows. Here, for example, are the χ^2 curves for 5 and 9 degrees of freedom.

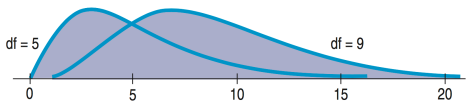

Notice that the value $\chi^2 = 10$ might seem somewhat extreme when there are 5 degrees of freedom, but appears to be rather ordinary for 9 degrees of freedom. Here are two simple facts to help you think about χ^2 models:

- The mode is at $\chi^2 = df - 2$. (Look back at the curves; their peaks are at 3 and 7, see?)
- The expected value (mean) of a χ^2 model is its number of degrees of freedom. That's a bit to the right of the mode—as we would expect for a skewed distribution.

Our test for zodiac birthdays had 11 df, so the relevant χ^2 curve peaks at 9 and has a mean of 11. Knowing that, we might have easily guessed that the calculated χ^2 value of 5.094 wasn't going to be significant.