[6] If the observed counts don't match the expected, the χ^2 statistic will be large. It can't be "too small." That would mean that our model really fit the data well. So the chi-square test is always one-sided. If the calculated statistic value is large enough, we'll reject the null hypothesis.

Even though its mechanics work like a one-sided test, the interpretation of a chi-square test is in some sense many-sided. With more than two proportions (Success/Failure), there are many ways the null hypothesis can be wrong. By squaring the differences, we made all the deviations positive, whether our observed counts were higher or lower than expected. There's no direction to the rejection of the null model. All we know is that it doesn't fit.

Some other characteristics of χ^2 models:

- Curves in the χ^2 family change both shape and center as the number of degrees of freedom grows.
- Unlike Normal and t families, χ^2 models are skewed right and cannot take on negative values.
- Unlike Normal and t-models, χ^2 models do not measure distances in standard deviation (or standard error) units.

With *z* or *t* a calculated result of 5 is pretty clearly significant; we can tell even without knowing the P-value, as it's rare for anything to be five standard deviations from what is anticipated. With the chi-square statistic the picture is murkier. Values that are highly significant when there are a few degrees of freedom fade into non-significance as the curve stretches to the right for greater degrees of freedom.

a) Since *df = # of categories* – 1, larger df are accompanied by more cells in the summation of the χ^2 statistic. Thus it takes a much larger value to surprise us when there are more degrees of freedom.
b) Look down the column of the χ^2 table to see how the critical value increases with higher degrees of freedom.
c) Analyzing various χ^2 distributions reveals: Mode = *df* – 2, expected value (mean) = *df*.
Having that initial feel for whether or not the calculated value is unusually large helps make the P-value a little less mysterious. The tables and distributions show a *rough* guide: unusually large = *df* * 2 (for *df*<15)

To calculate the exact P-value we just use our calculator: χ^2cdf(χ^2, 999, *df*)
We use 999 because unlike *t* and *z*, chi-square values can get big, especially when there are many cells.