

M&Ms!

introduction to randomness and hypothesis testing

Candy!!!

- Let's all get some candy from the bag, and then find the count of blue and count of white. Then we'll log it in google sheets.
- Next, let's use a function to find the proportion of blue for each person.
- What do you notice about the values we get?

The central limit theorem

The **central limit theorem** states that the sampling distribution of the mean of any **independent, random variable** will be normal or nearly normal, if the sample size is large enough.

In practical terms, this means that if we take enough samples, and they're big enough, the proportions will be normally distributed around the actual proportion in the population - in this case our whole bag of candy.

The central limit theorem

The central limit theorem lets us do a number of tests to see if our result is a fluke, or if it should cause us to question our assumptions.

Expectations

Often we have a belief about a characteristic of our data. This is called the 'null hypothesis.'

When we look at data, however, we might get a different answer. This is called our 'alternate hypothesis.'

The question then becomes, at what point can we accept our alternate hypothesis and reject the expected null hypothesis.

Hypothesis

We are going to test whether the actual proportion of blue M&Ms matches our expectation.

Because we are dealing with proportions, there are generally three null hypotheses.

- a. Proportion of blue = Proportion of white
- b. Proportion of blue $>$ Proportion of white
- c. Proportion of blue $<$ Proportion of white

Which one seems the most reasonable here?

Testing our hypothesis

There are many tests for many different questions. Because our hypothesis is comparing proportions, we will use what is called a z-test.

You don't need to remember what test gets used when. The important thing is to understand the concept of hypothesis testing, and know how to look for the right test.

Testing our hypothesis

The steps you should always use:

1. Prepare your data
2. Test the conditions and assumptions specific to the test you want to execute
3. Define your null and alternate hypotheses
4. Run the test

Test conditions and assumptions

- The sampling method for each population is simple random sampling.
- The samples are independent.

Testing our hypothesis

We meet all of the conditions. So let's go ahead.

Google sheets has limited capacity for statistical calculations (at least for now).

Luckily there are many online calculators.

Let's go here: <http://bit.ly/1PnSaNx>

Using the online calculator

Here's our result.

The key idea here is the p-value. For 95% significance level (which you will almost always use) you want a p-value that is less than .05.

Do we get that here?

z-statistic	2.000
Significance level	P = 0.0455
95% CI of observed proportion	30.33 to 50.28

What test do I use?

You don't need to know off the top of your head, but here are some examples:

