

# Lab 6 - Math 58 / 58b: power, 2 proportion z-tests

*your name here*

*due Mar 3, 2020*

## Lab Goals

- Using z-tests (normality) to calculate power
- How big a sample size for what difference in proportions with what baseline probability?

That is, in this lab, you will calculate the power to reject the null hypothesis by changing:

- (a) the sample size
  - (b) the difference in proportions (called the “effect size”)
  - (c) the baseline proportion
- how does the power depend on each of the three variables above?

## Getting started

Below we will calculate a series of values, so we will not actually do any data analysis. But let’s consider a two sample scenario just to have language to set up the problem better.

Consider a set-up where a placebo-controlled randomized trial proposes to assess the effectiveness of an antibiotic in reducing sepsis (infection) in premature babies. A previous study has shown the underlying rate of sepsis to be about 50% in such infants around 2 weeks after birth, and a reduction of this rate to 34% would be of clinical importance. The experiment will be conducted by giving half of the babies a placebo and half of the babies an antibiotic; whether or not they show signs of sepsis is recorded as the response variable. The test is one-sided, as we are only interested in the situation where the antibiotic reduces the rate of sepsis.

## Load packages

All of the work will be done using the normal approximation, so you will only need `xpnorm` and `xqnorm` in the `mosaic` package. Note that you can use `plot=FALSE` once you understand what the function reports (to make your lab write-up take fewer pages).

```
library(tidyverse)
library(mosaic)
```

## The initial approach

As mentioned above, we hope to see a difference in reduction of sepsis rates from 0.5 to 0.34 after using a placebo. Before running the study, use 0.5 and 0.34 (and not the pooled estimate from the data – heck, there is no data in this problem!) to find the power of the test when the experiment is set up so that there are 20 infants who get each drug.

1. First find the values of  $\hat{p}_1 - \hat{p}_2$  that will lead to rejecting  $H_0$ .
2. Now consider the situation where the antibiotic treatment results in only 34% of babies having sepsis (that is,  $H_0$  is no longer true. Indeed, the effect size is now 0.16). How often would your sampling mechanism produce a  $\hat{p}_1 - \hat{p}_2$  that would reject  $H_0$ ?
3. Sketch by hand (with a pencil, you won’t turn this in), the two normal curves from #1 and #2 above. Draw them to scale so that the power calculated in #2 is reasonably accurate on the sketch. [Note: drawing such a picture with a pencil is exactly the type of thing that could be asked on an in-class exam.]

## To turn in

4. Come up with list of possible values for:

- the sample size (try at least 5 different values)
- the difference in proportions (try at least 3 different values)
- the baseline proportion (try at least 2 different values)

[You may want to revisit / report later your choices for #4 after you've done #5.]

5. Find the power of the test for all combinations of the variables above. Most of the code has been written for you as three different nested for loops. The three vectors `samp_sizes`, `diff_props`, and `base_props` will need to be filled in with the values from above.

```
samp_sizes <- c() # add per group sample size here
diff_props <- c() # add the effect size here
base_props <- c() # add the baseline probability (placebo)
```

Don't change any of the code but **do** read through the code to make sure you understand what it is doing:

```
test_samp_size <- c()
test_diff_props <- c()
test_base_props <- c()
test_power <- c()

for(n in samp_sizes){
  for(d in diff_props){
    for(b in base_props){
      a = b - d # antibiotic proportion
      cutoff <- xqnorm(0.95, 0, sqrt(b*(1-b)/n + b*(1-b)/n),
                     verbose= FALSE, plot = FALSE)
      power <- 1 - xpnorm(cutoff, d, sqrt(b*(1-b)/n + a*(1-a)/n),
                         verbose = FALSE, plot = FALSE)

      test_samp_size <- c(test_samp_size, n)
      test_diff_props <- c(test_diff_props, d)
      test_base_props <- c(test_base_props, b)
      test_power <- c(test_power, power)
    } } }

power_results <- data.frame(test_power, test_samp_size,
                             test_diff_props = as.factor(test_diff_props),
                             test_base_props = as.factor(test_base_props))
```

6. Look at your results. How many rows does it have? What does each row represent? How many columns? What does each column represent? Print a bit of your data (not all of it!).

7. Using `ggplot` create plots with the following details:

- `test_power` on the y-axis (use a line geom)
- `test_samp_sizes` on the x-axis
- group by the `test_diff_props` variable (note that above I transformed the variable from a number into a character)
- color by the `test_diff_props` variable
- add one additional layer to “facet” by the variable `test_base_props`. The last layer of your plot will be: `facet_wrap(~test_base_props)`.

8. What does the plot tell you about what you should ask for in terms of funding? That is, how many observations in each group do you need? (Hint: what does “need” mean?)