

19.3: Simulating Statistical Power

Let's simulate this to see whether the power analysis actually gives the right answer. We will sample data for two groups, with a difference of 0.5 standard deviations between their underlying distributions, and we will look at how often we reject the null hypothesis.

```
nRuns <- 5000
0.5effectSize <-
# perform power analysis to get sample size
pwr.result <- pwr.t.test(d=effectSize, power=.8)
# round up from estimated sample size
sampleSize <- ceiling(pwr.result$n)

# create a function that will generate samples and test for
# a difference between groups using a two-sample t-test

get_t_result <- function(sampleSize, effectSize){
  # take sample for the first group from N(0, 1)
  group1 <- rnorm(sampleSize)
  group2 <- rnorm(sampleSize, mean=effectSize)
  ttest.result <- t.test(group1, group2)
  return(tibble(pvalue=ttest.result$p.value))
}

index_df <- tibble(id=seq(nRuns)) %>%
  group_by(id)

power_sim_results <- index_df %>%
  do(get_t_result(sampleSize, effectSize))

p_reject <-
  power_sim_results %>%
  ungroup() %>%
  summarize(pvalue = mean(pvalue<.05)) %>%
  pull()

p_reject
```

```
## [1] 0.8
```

This should return a number very close to 0.8.

This page titled [19.3: Simulating Statistical Power](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Russell A. Poldrack](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.