17.2: Simulating p-values

In this exercise we will perform hypothesis testing many times in order to test whether the p-values provided by our statistical test are valid. We will sample data from a normal distribution with a mean of zero, and for each sample perform a t-test to determine whether the mean is different from zero. We will then count how often we reject the null hypothesis; since we know that the true mean is zero, these are by definition Type I errors.

```
nRuns <- 5000

# create input data frame for do()
input_df <- tibble(id=seq(nRuns)) %>%
  group_by(id)

# create a function that will take a sample
# and perform a one-sample t-test

sample_ttest <- function(sampSize=32){
  tt.result <- t.test(rnorm(sampSize))
  return(tibble(pvalue=tt.result$p.value))
}

# perform simulations

sample_ttest_result <- input_df %>%
  do(sample_ttest())

p_error <-
  sample_ttest_result %>%
  ungroup() %>%
  summarize(p_error = mean(pvalue<.05)) %>%
  pull()

p_error
```

```
## [1] 0.048
```

We should see that the proportion of samples with $p < .$ is about 5%.