Power

The goal of this project is to understand the power function.

We will use the following functions. Copy and past the whole amount into the R command line:

```
sim = function(n = 10, realmu=0, pop="norm", m = 250)  {
  ## pop either "norm", "long", or "skew"

  res = c()

  if(pop =="norm") {
    x = matrix(rnorm(n*m,realmu,1),ncol=m)
  } else if (pop == "long") {
    x = matrix(rt(n*m,df=3)+realmu, ncol=m)
  } else {
    x = matrix(rexp(n*m,1)+realmu - 1, ncol=m)
  }

  ## apply function to each column
  res = apply(x,2,function(x) t.test(x)$p.value)


  return(res)

}

## What proportion?
prop = function(x) sum(x)/length(x)
## find CI from logical vector
Prop.test = function(x, ...) prop.test(sum(x), length(x),...)
```

The function `sim()` computes p-values for a simulation with

$$H_0 : \mu = 0 \qquad H_A : \mu \neq 0$$

using a t-test.

The simulation produces m p-values for samples of size n (default 10) from a population with true mean `realmu` with a default of 0 (H_0 is true). The population by default is normal, but setting `pop="t"` will use a long tailed distribution and `pop="exp"` will use a skewed distribution.

1 Finding α and β through simulation

The probability of a Type-I error is called α. This is the probability that we reject H_o when it is indeed true. To see how we can simulate α, run `sim()` without any arguments:

```
> res = sim()
> prop(res < 0.05)

[1] 0.04
```

When a p-value is less than 0.05 we "reject" in this case. The proportion should be close to 0.05.

Question 1: Why should the proportion be close to 0.05?

The `prop.test()` function would calculate a CI for the value of α. We use `Prop.test()` so that we can use unsummarized data:

```
> Prop.test(res < 0.05)

        1-sample proportions test with continuity correction
data:  sum(x) out of length(x), null probability 0.5
X-squared = 209.764, df = 1, p-value < 2.2e-16
alternative hypothesis: true p is not equal to 0.5
95 percent confidence interval:
 0.02045396 0.07456195
sample estimates:
   p
0.04
```

Question 2: Does this CI contain 0.05?

Question 3: Why is it that the simulation of the t statistic is analyzed using a test of proportions?

Question 4: Repeat the above, only use `sim() < 0.10`. What is α now? Can you conjecture what it would be for `sim() < 0.03`?

1.1 Finding α for different populations

The above simulations used a normal population. What if we use a long tailed distribution (the t with 3 degrees of freedom)?

```
> Prop.test(sim(pop = "t") < 0.05)

        1-sample proportions test with continuity correction
data:  sum(x) out of length(x), null probability 0.5
X-squared = 174.724, df = 1, p-value < 2.2e-16
```

```
alternative hypothesis: true p is not equal to 0.5
95 percent confidence interval:
 0.05078401 0.12266504
sample estimates:
   p
0.08
```

Question 5: Does the CI contain 0.05?

Question 6: Repeat the above with the arguments

```
> sim(n = 4, pop = "t")
```

and

```
> sim(n = 100, pop = "t")
```

Is the result the same for each?

Question 7: For $n = 4, 10$, and 100 find a CI for α with `pop="exp"` and 0.05 in the inequality.

1.2 Finding β for different μ and n

The function $\beta(\mu)$ computes the probabilty of accepting the null hypothesis given a value of the parameter μ. If μ is different than that assumed in H_0, then this is when H_0 is false.

The power function is $1 - \beta(\mu)$. This means small βs are bigger.

To see how often we falsely accept when $\mu = 1$ (not 0), $n = 10$ and $\alpha = 0.05$ we have

```
> alpha = 0.05
> Prop.test(sim(n = 10, realmu = 1) > alpha)

        1-sample proportions test with continuity correction
data:  sum(x) out of length(x), null probability 0.5
X-squared = 81.796, df = 1, p-value < 2.2e-16
alternative hypothesis: true p is not equal to 0.5
95 percent confidence interval:
 0.1640908 0.2689662
sample estimates:
    p
0.212
```

Question 8: What is the 95% CI for $\beta(1)$ when $n = 10$?

Question 9: What is the 95% CI for $\beta(1)$ when $n = 100$?

Question 10: Is β less than 0.10 when $\mu = 1$ and $n = 25$?

Question 11: For $n = 10$ find a CI for β when $\mu = .1, .5, 1, 2$.

Question 12: For $n = 100$ find a CI for β when $\mu = .1, .5, 1, 2$.

Question 13: Comment on the value of $\beta(\mu)$ as mu increases from 0 for a fixed n? (What is the value at $\mu = 0$. (Assume $\alpha = 0.05$)

Question 14: Comment on the value of $\beta(\mu)$ for a fixed mu as n gets big. (Assume $\alpha = 0.05$)

2 Is the t test more powerful?

We know that for simple hypotheses the likelihood ration test is most powerful, but what about which of these is stronger: the t-test or the signed-rank test?

Fix $n = 10$ and $\mu = 1$. Then define the data by

```
> m = 250; n = 10
> x = matrix(rnorm(m * n, mean = 1, sd = 1), ncol = m)
> res1 = apply(x, 2, function(x) wilcox.test(x)$p.value)
> res2 = apply(x, 2, function(x) t.test(x)$p.value)
```

Question 15: Which test is more powerful for this data? How do you know?

3 The function `power.t.test()`

Power calculations are built into R. The function `power.t.test()` will do the work for us. For instance, to find the power $(1 - \beta)$ when $n = 10$ and the difference in means (real μ minus that assumed in H_0) is 1 we haev

```
> power.t.test(n = 10, delta = 1)

     Two-sample t test power calculation
              n = 10
          delta = 1
             sd = 1
      sig.level = 0.05
          power = 0.5619846
    alternative = two.sided
 NOTE: n is number in *each* group
```

Question 16: Is the power more than 0.90? If not, try different values of n to find out how large n should be to get a power greater than 0.90.

Question 17: That last exercise was a tad tedious, verify that the command

```
> power.t.test(delta = 1, power = 0.9)
```

will do that work more quickly.