

**3. Modify Fisher's Sharp Null to reflect the null hypothesis that the treatment adds five units to the outcome under control. Build a small simulation (e.g., 10 observations) and construct a table that summarizes the potential outcomes. Randomize using a fair coin flip to assign treatment or control for each observational unit. Use the permutation test to assess your data set using (i) Fisher's Sharp Null and (ii) the null hypothesis that the treatment adds five units to the outcome under control.**

The first part is straightforward enough: Fisher's Sharp Null is  $H_0: r_{C_i} = r_{T_i}$ . If we instead believe that the treatment adds "five units" to the response that would appear under the control then we're saying we believe  $r_{C_i} + 5 = r_{T_i}$  which can be made into a null hypothesis by slapping this notation on it:  $H_0: r_{C_i} + 5 = r_{T_i}$ .

The next part of the question asks us to build a simulation... but gives us basically nothing to work with in how to set this up. You thought this was a stats class and now you find yourself in the middle of a choose-your-own-adventure-book. The heck is this? Chill. Inventory what you have, inventory what you don't have.

If we keep reading the rest of the question we see that we'll be randomizing with a fair coin. And we're going to use a permutation test. Looks like we need a  $r_{C_i}$ ,  $r_{T_i}$  and  $Z_i$  for all  $i \in [0,10]$ . We don't have any knowledge about what  $r_{C_i}$  looks like. We don't even have dimensions. But we do have something about  $r_{C_i} + 5 = r_{T_i}$ . Is that how they're actually connected? Nope. That's a hypothesis, a conjecture about the world. The actual relationship is not declared. We don't see anything about how to connect the potential outcomes... so let's just choose something<sup>1</sup>.

Let's generate the  $r_{C_i}$  by drawing from a  $N(0,10)$  for each  $i$ . Why'd I choose sd=10? No real reason, I kinda chose it to be big because it seems like we're curious if there's some treatment effect on the order of "5 units," so I figure if the sd was small then it'd be really easy to tell if the treatment existed. Thus I decided to make the sd big because maybe things will be interesting that way. Why normal? Why not? (I did think about using a uniform distribution...)

More interesting, is now we have to decide how to generate  $r_{T_i}$ . If we want, we could set  $r_{C_i} = r_{T_i}$ . Or we could generate  $r_{T_i}$  completely independent from  $r_{C_i}$ . There's nothing stopping us from doing that. Perfectly valid. We're playing the role of Nature right now; our motivations can be as mysterious as we'd like. I'm going to set  $r_{C_i} + 2 = r_{T_i}$ . (But I could set it to  $2r_{C_i} = r_{T_i}$  and this would all be ok.)

Next we'll generate all of the  $Z_i$ . They'll come from Bernoulli( $p=0.5$ ).

---

<sup>1</sup> This is our way of emphasizing that the permutation test doesn't make any distributional assumptions. The covariates can behave whatever way they please; we just need to know how the  $Z_i$  are generated.

Turning to R, I'll execute the following code to generate the data:

```
#install.packages("prob") #install package if you haven't already done so
library(prob) #load the package
set.seed(432) # specific seed to reproduce numbers below
n <- 10 #10 observations
r_c <- rnorm(n,mean=0,sd=10) #r_c are generated from N(0,10)
r_t <- r_c + 2 #treatment adds 2 to r_c
z <- rbinom(n,size=1,prob=0.5) #n fair coin flips
r_obs <- z*r_t + (1-z)*r_c #get the observed outcomes
delta <- r_t - r_c #individual level effects
cbind(1:10,r_c,r_t,z,r_obs,delta) #potential outcomes table
```

Here's the potential outcomes table:

person	r_c	r_t	z	r_obs	delta
1	-6.6	-4.6	0	-6.6	2
2	14.0	16.0	1	16.0	2
3	6.0	8.0	1	8.0	2
4	6.3	8.3	1	8.3	2
5	4.7	6.7	0	4.7	2
6	3.9	5.9	0	3.9	2
7	-6.7	-4.7	0	-6.7	2
8	-4.8	-2.8	0	-4.8	2
9	-0.8	1.2	1	1.2	2
10	-1.9	0.1	1	0.1	2

If we weren't pretending to be omniscient then the data table would look more like this:

person	z	r_obs
1	0	-6.6
2	1	16.0
3	1	8.0
4	1	8.3
5	0	4.7
6	0	3.9
7	0	-6.7
8	0	-4.8
9	1	1.2
10	1	0.1

From the observed table immediately above, we can't recreate the potential outcomes table.

From the observed table above, we can only recreate this much of the potential outcomes table:

person	r_c	r_t	z	r_obs	delta
1	-6.6		0	-6.6	
2		16.0	1	16.0	
3		8.0	1	8.0	
4		8.3	1	8.3	
5	4.7		0	4.7	
6	3.9		0	3.9	
7	-6.7		0	-6.7	
8	-4.8		0	-4.8	
9		1.2	1	1.2	
10		0.1	1	0.1	

If we assume Fisher's Sharp Null hypothesis is true then we can put in values for the delta, and that implies what the unobserved values must be.

person	r_c	r_t	z	r_obs	delta
1	-6.6	-6.6	0	-6.6	0
2	16.0	16.0	1	16.0	0
3	8.0	8.0	1	8.0	0
4	8.3	8.3	1	8.3	0
5	4.7	4.7	0	4.7	0
6	3.9	3.9	0	3.9	0
7	-6.7	-6.7	0	-6.7	0
8	-4.8	-4.8	0	-4.8	0
9	1.2	1.2	1	1.2	0
10	0.1	0.1	1	0.1	0

**Permutation test under Fisher's Sharp Null:** To do the permutation test we need to decide what our test statistics is. We can choose anything we want. A very common one is to take the average of the responses for the treated group and subtract the average of the responses for the control group, that is  $mean(r_{obs}|z = 1) - mean(r_{obs}|z = 0)$ . In this example we have an observed test statistics of approximately 8.62. The permutation test is pretty simple under Fisher's Sharp Null, we basically just flip whether the observed value came from the treated or control group. Personally, I tend to do only one-sided tests because the alternative in most situations is one-sided (e.g., I'm testing this drug because I believe it's going to help.)

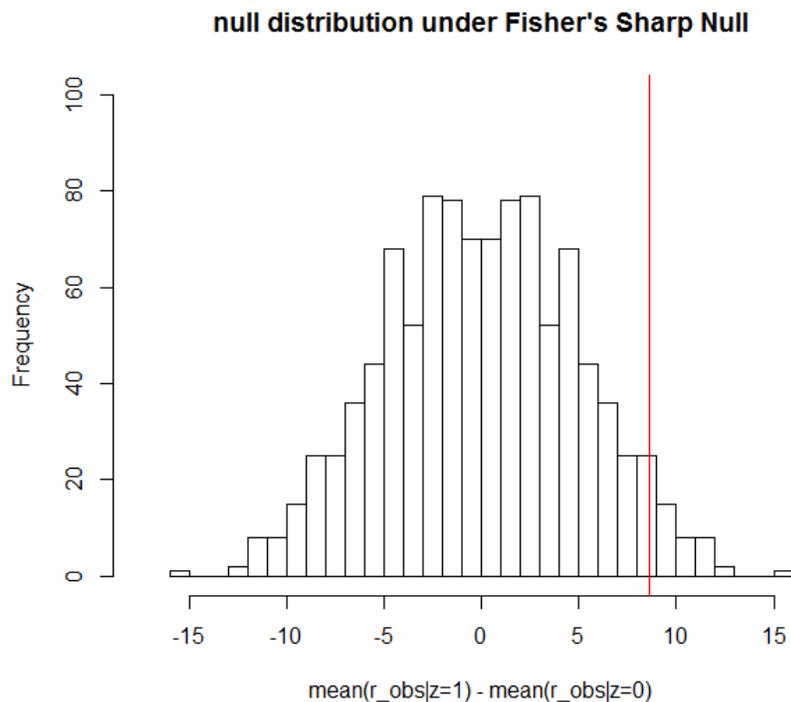
Note that two of the extremes – i.e., z comes up all heads or all tails) – present us with a problem. What would you do with them? I'd argue that they are uninformative of the treatment effect and thus are excluded from the space of all acceptable random assignments.

Here's the R code I used to turn the permutation test under Fisher's:

```
install.packages("prob")
library(prob) #load the package
sample.space <- 1*(tosscoin(10)== "H")

#####
##Fisher's Sharp Null##
#####
bar_r_t <- (sample.space%%r_obs)/apply(sample.space,1,sum)
bar_r_c <- ((1-sample.space)%%r_obs)/apply((1-sample.space),1,sum)
diff <- (bar_r_t-bar_r_c)
hist(diff,
      breaks="fd",
      main="null distribution under Fisher's Sharp Null",
      xlab="mean(r_obs|z=1) - mean(r_obs|z=0)")

obs_diff <- sum(z*r_t)/sum(z)-sum((1-z)*r_c)/sum(z) #get the observed diff
abline(v=obs_diff,col="red") #add a red line to show the observed diff
p_value <- sum(diff>obs_diff,na.rm=TRUE)/(length(diff)-2) #there are two NAs
```



The p-value calculates the portion of the null distribution that is greater than the observed statistic. In this case we get a p-value of about 0.038.

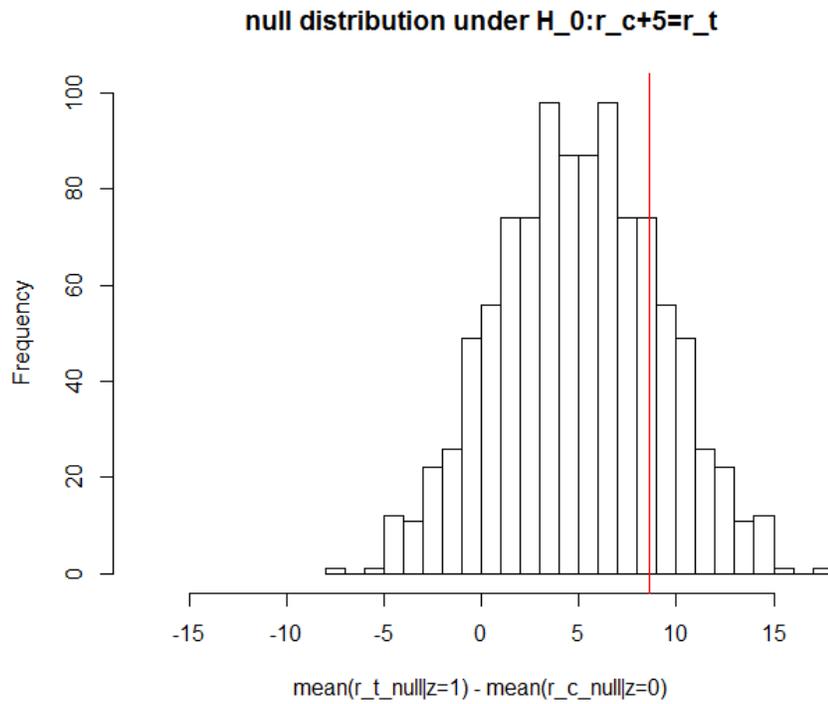
**Permutation test under  $H_0: r_{C_i} + 5 = r_{T_i}$ :** This is a bit more complicated than under Fisher's. Now we need to use the null hypothesis to get the unobserved counterfactuals. From the null we have  $\Delta_i = r_{T_i} - r_{C_i} = 5$ . Filling in the potential outcomes table we get:

person	r_c	r_t	z	r_obs	delta
1	-6.6	-1.6	0	-6.6	5
2	11.0	16.0	1	16.0	5
3	3.0	8.0	1	8.0	5
4	3.3	8.3	1	8.3	5
5	4.7	9.7	0	4.7	5
6	3.9	8.9	0	3.9	5
7	-6.7	-1.7	0	-6.7	5
8	-4.8	0.2	0	-4.8	5
9	-3.8	1.2	1	1.2	5
10	-4.9	0.1	1	0.1	5

This means we need to adjust our code a bit:

```
#####
##H_0:r_c+5=r_t##
#####
r_t_null <- ifelse(z==1,r_obs,r_obs+5)
r_c_null <- ifelse(z==0,r_obs,r_obs-5)
bar_r_t <- (sample.space%%r_t_null)/apply(sample.space,1,sum)
bar_r_c <- ((1-sample.space)%%r_c_null)/apply((1-sample.space),1,sum)
diff <- (bar_r_t-bar_r_c)
hist(diff,
      breaks="fd",
      main="null distribution under H_0:r_c+5=r_t",
      xlab="mean(r_t_null|z=1) - mean(r_c_null|z=0)")

obs_diff <- sum(z*r_t_null)/sum(z)-sum((1-z)*r_c_null)/sum(z)
abline(v=obs_diff,col="red")
p_value <- sum(diff>obs_diff,na.rm=TRUE)/(length(diff)-2)
```



In this case, if we do a one-sided test for greater than, we get a p-value of approximately 0.199.<sup>2</sup>

---

<sup>2</sup> Unlike the test of “no effect” earlier, you might reasonably go into this test considering a two-sided alternative hypothesis. In that case the p-value would be approximately  $2*(0.199) = 0.398$ .