

Handed out: Wednesday January 20, 2010

Due: Wednesday February 3, 2010

**NOTE:** There is some computing required for this assignment. You can use any software or programming language that you like. A sample R program that does the exact (and simulated) p-value calculation for the  $n = 6$  honey experiment in Chapter 5 is provided at the end of the assignment. Please email or come see me if you need help.

1. **Completely randomized study:** The Table below gives data from a small study analyzing the effect of Vitamin A supplements on the immune system's capability to fight malaria. Twelve children participated. Six were randomly assigned to receive Vitamin A supplements and the other six were assigned to a control treatment (placebo). The outcome is a measurement of the quantity of malaria parasites presents (the parasite load measured on the log scale).

Placebo	Vitamin A
8.62	0.06
1.48	1.72
8.93	2.19
9.57	7.32
2.65	7.53
7.30	7.62

Suppose we want to test Fisher's (sharp) null hypothesis of no treatment effect on these 12 children. Take the test statistic to be the difference in sample means.

- Calculate the exact p-value (two-tailed). Hint: You will need to write a program to find the 924 possible values of the test statistic; see the sample program below.
  - Consider a simulation approximation to the p-value: generate 1000 simulated assignments, compute the test statistic for each, and determine the proportion of trials with a larger outcome than the observed statistic. Again, see the sample program below.
  - How accurate is the simulation approximation? Explain by referring to the expected variability in a sample proportion of this type.
  - Report the p-value you would obtain using a traditional t-test to compare the two group means (on the log scale). Is this related to the p-value in (a)? Explain.
  - Obtain a (crude) confidence interval for the treatment effect by assuming  $Y_i(1) = Y_i(0) + k$  and then testing this assumption for different values of  $k$ . Hint: Use a grid of  $k$ -values and then compute/simulate as in (a) or (b).
2. **Randomized study with pairing.** We now consider the same data as in (1) under a different assumption. Suppose that prior to the experiment, researchers used a characteristic  $X$  thought to be related to the outcome to match the children in 6 pairs (the 6 pairs being the rows of Table 1) and then randomly assigned the treatment to one unit within each pair.
- Explain how this changes the execution of Fisher's test and obtain the exact p-value for the new test.
  - Carry out a simulation to get a Monte Carlo approximation to the p-value.
  - Report the p-value you would obtain using a traditional t-test in this case (it would be the paired t-test).
  - Discuss why randomization within pairs might be a good design even though it restricts the set of allowable randomizations. Tell whether it appears to have been a good idea here.
3. **Theory.** We use the same data here for a short theory problem.
- Write the test statistic (the difference in sample means) as an explicit function of the potential outcomes and the assignment indicators for the 12 units.
  - Show that under complete randomization (as in (1)) the statistic is unbiased for the average treatment effect in the 12 units.
  - Is it still unbiased for the pairwise randomization? Explain.

4. **Project.** A major evaluation in the class will be a project. You can work alone or in a group of two. There is some flexibility here in that you can: (1) focus on a topic that we don't cover (see below); (2) review the literature in a controversial area where causal arguments play a role; or (3) carry out an analysis (or reanalysis) of data from an observational study using methods in the class. For this assignment your responsibility is to identify your partner (if any), list one (or more) possible topics, and write a paragraph on what you are thinking about for the topic(s). (I say one or more topics in case you want feedback on more than one idea).

In terms of methodology – possible topics include regression discontinuity designs, sensitivity analysis / bounds for causal effects, methods for sequential treatment problems (marginal structure models), Pearl's approach to causal inference, etc. For application areas virtually anything is possible including effects of global warming, SAT coaching, education vouchers, etc. There are multiple articles about some or all of these topics in the reading lists that I posted on the website with this homework. For your own analysis – you/we would need to find a data set like the ones used in Imbens/Rubin or elsewhere.

### Sample R programs for (1) and (2)

```
#
# PROGRAM TO COMPUTE RANDOMIZATION DISTRIBUTION FOR HONEY DATA IN CHAPTER 5
#
# set data with "nt" treatment obs followed by "nt" control obs
n <- 6
yobs <- c(3,5,0,4,0,1)
nt <- n/2
#
# compute observed test statistic
sumy <- sum(yobs)
tobs <- sum(yobs[1:nt])/nt - sum(yobs[(nt+1):n])/nt
#
# set up vector to hold randomization distribution of test stat
out <- rep(0,choose(n,nt))
cnt <- 0
#
# loop to identify unique randomizations
for (i1 in (1:(nt+1))) {
  for (i2 in ((i1+1):(nt+2))) {
    for (i3 in ((i2+1):n)) {
#
# compute test statistic for randomization and store
      cnt <- cnt + 1
      trtsum <- sum(yobs[c(i1,i2,i3)])
      out[cnt] <- trtsum/nt - (sumy-trtsum)/nt
    }}}
#
# compute p-value
pval <- sum(abs(out) >= abs(tobs))/choose(n,nt)

#
# SIMULATION VERSION
# uses the R sample command to choose the sample
outsim <- rep(0,1000)
for (i in (1:1000)) {
  ysamp <- sample(yobs,nt)
  trtsum <- sum(ysamp)
  outsim[i] <- trtsum/nt - (sumy-trtsum)/nt
}
pvalsim <- sum(abs(outsim) >= abs(tobs))/1000
```