

## Pre-Class Response for Lecture 5:

Suppose you are running an RCT to decide if a new drug is effective. People enter the study one at a time, say, about once a week, as they come into the hospital with symptoms. You randomly assign each study participant to either the treatment or control group, and you can tell within a week if the patient recovers. You intend to include 500 people in the study.

**Question 1:** Suppose that, halfway through the study, after you have seen 250 people, you realize that everyone in the treatment group (about 125 people) got better, and everyone in the control group (about 125 people) did *not* get better. What should you do? Should you complete the study? Why or why not? What if about 85% of the treatment group got better and 85% of the control group did *not* get better?

**Question 2:** This situation differs from the sort of RCT that we discussed in class, in that it could be "*online*" (also called "*adaptive*"). That is, people enter the study one at a time, and you can see the outcome before the next person enters, so **your choice of whether to assign the  $n$ 'th person to the treatment or control group can depend on what happened to people 1, 2, ...,  $n-1$ .** In contrast, in the sort of RCT we discussed earlier, you assign everyone to treatment or control at the beginning, and then see the outcome for everyone at the end. If your goal is to help as many people as possible to recover (both the people in the study and all the people to come), can you take advantage of this "online"-ness to do better than a standard RCT? If so, what would you do? And if not, why not? Answer in a paragraph or two. (In particular, if your answer is "yes," we aren't asking for a fully specified protocol and a proof that it works, just some vibes about what you might do and why it seems reasonable).