Axioms and assumptions

18.600 Problem Set 2, due February 19

Welcome to your second 18.600 problem set! We are still thinking about combinatorics and probability, and about the axioms of probability.

Before we get to work, let's indulge in just a bit of reflection. When we say "The probability that A will happen is p" where does p come from? Sometimes the scientific evidence is strong enough to convince pretty much everyone that A will or will not happen. Informally, we might say that the probability that a predicted lunar eclipse will happen on schedule is pretty much 1, and the probability that Mars and Jupiter will collide this month is pretty much 0. In other simple situations (die rolls, coin tosses, etc.) experience may lead us to agree on probabilities that aren't 0 or 1. The assumption that all outcomes are equally likely (for random permuations or die rolls or coin tosses) is sometimes a natural starting point. This assumption is implicitly made in a few of the problems on this problem set.

In other real world settings, one can define the *risk neutral* probability, a probability measure derived from the market prices of contracts whose values depend on future events. If we want to know the risk neutral probability that a given candidate will win an election, or that an athletic team will win a game, we can look at betting markets. (Check out predictwise.com and similar websites.) As we will see later in the course, if we want to know the risk neutral probability that the price of a share of Apple stock will exceed some value by the end of the year, we can work this out by looking at current prices of *derivatives* (contracts whose future value depends on future share prices). The total amount of money at stake in derivative markets is estimated at over a quadrillion dollars per year (try googling derivatives quadrillion).

Some argue that betting markets set up perverse incentives. If I buy a contract that gives me $500,000$ dollars if my house burns down, that's a useful form of insurance. But if I buy a contract that gives me $500,000$ dollars if *your* house burns down, that's more problematic: it gives me an unhealthy incentive to burn your house down. People similarly worry about a world in which hedge funds can bet that a company will collapse and then actively cause it to collapse. Rules are required to prevent such things.

On the other hand, one might argue that the absence of betting markets is part of the reason that some questions in politics and law are divisive. It is hard to place a bet on the proposition that "my candidate would do more to advance long term happiness and prosperity than yours" or "my client is innocent," so there is no market mechanism for producing a commonly accepted probability. Different groups can *claim* to have different probability estimates, the expression of which may advance their own agendas, but without a market we cannot tell which parties would actually be *willing to bet money* at the corresponding rates. Some studies claim that people answering questions about the economy are both more accurate and less partisan when they are paid (even a very small amount) for correct answers. Maybe there is something to be said for having money on the line.

Legal systems around the world recognize many different "burdens of proof" including probable cause, reasonable suspicion, reasonable doubt, preponderance of evidence, beyond a shadow of a doubt, clear and convincing evidence, some credible evidence, and reasonable to believe. Most of these do not have a clear meaning as numerical probabilities (does "beyond reasonable doubt" mean with probability at least .95, or at least .99, or something else?) but there is an exception: "preponderance of evidence" is generally taken to indicate that a probability is greater than fifty percent, so that something can be said to be "more likely than not." An interesting question (which I am not qualified to answer) is whether numerical probabilities should be assigned to the other terms as well. As a fanciful proposal for a new jury system, imagine telling a jury that you (might) secretly know whether the defendant is guilty, then asking jurors to form a betting market to bet on whether the defendant is guilty, then checking to see if the market prices reflect a guilt probability of greater than .95, and then convicting the defendant if they do. We could even ask the jurors to bet on a range of scenarios, so that we end up with a probability measure on a sample space with several elements, consistent with the axioms of probability.

Not very realistic, I know. But as long as we have *some* way to define probability measures for which the axioms of probability hold, we can get back to studying these measures with mathematics. So let's get to it.

Please stop by my weekly office hours (2-249, Wednesday 3 to 5) for discussion.

A. FROM ROSS 8TH EDITION CHAPTER TWO:

1. **Problem 25:** A pair of dice is rolled until a sum of either 5 or 7 appears. Find the probability that a 5 occurs first. *Hint.* Let E_n denote the event that a 5 occurs on the nth roll and no 5 or 7 occurs on the first $(n-1)$ rolls. Compute $P(E_n)$ and argue that $\sum_{n=1}^{\infty} P(E_n)$ is the desired probability.

2. **Problem 48:** Given 20 people, what is the probability that, among the 12 months in the year, there are 4 months containing exactly 2 birthdays and 4 containing exactly 3 birthdays?

3. **Theoretical Exercise 10:** Prove that $P(E \cup F \cup G) = P(E) + P(F) + P(G) - P(E^cFG) - P(EF^cG) - P(EFG^c) - 2P(EFG)$.

4. **Theoretical Exercise 20:** Consider an experiment whose sample space consists of a countably infinite number of points. Show that not all points can be equally likely. Can all points have a positive probability of occurring?

B. Suppose that there are 10 job candidates and 10 companies with job openings. Each candidate (independently, uniformly at random) develops a serious interest in one of 10 companies and each company (independently, uniformly at random) develops a serious interest in one of the 10 candidates. What is the probability that there is at least one company-candidate pair that are seriously interested in each other? (Hint: let E_j be the event that the jth applicant's interest is requited. Use inclusion-exclusion on these events. You can write the probability as a sum. Don't worry about simplifying further.)

C. Alice and Bob are playing a game of tennis and have reached the game state called "deuce." From here the players keep playing points until one player's total exceeds the other player's total by 2, at which point the player ahead by 2 points is declared winner of the game. Suppose that Alice wins each point with probability p (independently of all previous points) and Bob wins each point with probability $q = (1-p)$ (independently of all previous points). Find the probability that Alice wins the game, as a function of p. (Hint: consider what happens over the course of the next *two* points. Either Alice wins both and the game is over, or Bob wins both and the game is over, or each player wins a point and the players are back where they started. Compute the probabilities of these three outcomes. Then apply the ideas from the first problem on this problem set.) Based on your answer, do you agree or disagree with the following statement? *If Alice is k times as likely as Bob to win a point, then Alice is k^2 times as likely as Bob to win the game if the current score is deuce.*

D. The online comic strip xkcd.com has a "random" button one can click to choose one of the previous $n \approx 1642$ strips. Assume that there are exactly 1642 numbered strips and that each time one clicks the "random" button, one gets the kth strip where k is chosen uniformly from $\{1, 2, \dots, n\}$. If one observes m strips in this way, what is the probability that one sees at least one strip more than once? (This is a variant of the birthday problem.) Give approximate numerical values for $m = 30$ and $m = 50$ and $m = 70$. (Hint: try going to wolframalpha.com and entering something like N[Prod[(1-k/1642), {k,0,29}]]. In Mathematica, N[expression] produces a numerical approximation for the expression.) Based on your answer, do you agree or disagree with the following statement? *The number of clicks required before you see the same strip twice is a random quantity whose median is about* 50*, and which lies between* 30 *and* 70 *about half the time*

E. A lazy professor is teaching a class with 10 first year students, 10 second year students, 10 third year students, and 10 fourth year students. At the end of the semester, the professor assigns grades randomly: a group of 20 students is chosen uniformly at random (from all possible size 20 groups) to be the students who get A grades. What is the probability that exactly 5 students from each year end up getting A's? What is the probability that *all* of the first year students get A's?

F. (Just for fun – not to hand in.) The following is a popular and rather instructive puzzle. A standard deck of 52 cards (26 red and 26 black) is shuffled so that all orderings are equally likely. We then play the following game: I begin turning the cards over one at a time so that you can see them. At some point (before I have turned over all 52 cards) you say "I'm ready!" At this point I turn over the next card and if the card is red, you receive one dollar; otherwise you receive nothing. You would like to design a strategy to maximize the probability that you will receive the dollar. How should you decide when to say "I'm ready"?