Chapter 1

Some basic probability theory

1.1 Lecture 1 (3/Oct): Appetizers

1. Measure the length of a long string coiled under a glass tabletop. You have an ordinary rigid ruler (longer than sides of the table).

2. N gentlemen check their hats in the lobby of the opera, but after the performance the hats are handed back at random. How many men, on average, get their own hat back?

3. The coins-on-dots problem: On the table before us are 10 dots, and in our pocket are 10 nickels. Prove the coins can be placed on the table (no two overlapping) in such a way that all the dots are covered.

4. Birthday Paradox. I just remind you of this: a class of 23 students has better than even odds of some common birthday. (Supposing birthdates are uniform on 365 possibilities.) The exact calculation is
$$\Pr(\text{some common birthday}) = 1 - \frac{365 \cdots 343}{365^{23}} \cong 0.507297$$
but a better way to understand this is that the number of ways this can happen is $\binom{k}{2}$ for k students; so long as these events don't start heavily overlapping, we can almost add their probabilities (which are each just $\frac{1}{365}$). We'll be more formal about the upper and lower bounds soon.

5. The envelope swap paradox: You're on a TV game show and the host offers you two identical-looking envelopes, each of which contains a check in your name from the TV network. You pick whichever envelope you like and take it, still unopened.

   Then the host explains: one of the checks is written for a sum of $\$N$ ($N > 0$), and the other is for $\$10N$. Now, he says, it's 50-50 whether you selected the small check or the big one. He'll give you a chance, if you like, to swap envelopes. It's a good idea for you to swap, he explains, because your expected net gain is (with $\$m$ representing the sum currently in hand):
$$E(\text{gain}) = (1/2)(10m - m) + (1/2)(m/10 - m) = (81/20)m$$
   How can this be?

6. Consider a certain society in which parents prefer female offspring. Can a couple increase their expected fraction of daughters by halting reproduction after the first?

   Let's just make explicit here that we are not using advanced medical technologies. That is to say, the couple can control whether they create a pregnancy, but no other property of the fetus.

Before moving on we note that this is closely related to a famous problem which we will return to: the gambler's ruin problem. A gambler starts with \$1 in his pocket and repeatedly risks \$1 on a fair coin toss, until he goes broke. He is very likely to go broke, right? Indeed if he sticks around indefinitely, he will go broke with probability 1. But that is exactly equivalent (when boys and girls are equiprobable) to the event of a sufficiently large family having an excess of girls over boys.

Why are our intuitions so opposite in these two cases? It has to do with the fact that we clearly internalize the finiteness of the family size, whereas we can easily imagine the gambler addicitively playing for an extraordinarily long time. So his high-probability doom impresses us. If we decide in advance that we will stop him after one million plays, whether or not he has stopped himself by that time, then his expected wealth at that time is equal to \$1, even though he has almost certainly lost his \$1 and gone broke; there's a small chance he has earned a lot of money.

7. Unbalancing lights: You're given an $n \times n$ grid of lightbulbs. For each bulb, at position (i, j), there is a switch b_{ij}; there is also a switch r_i on each row and a switch c_j on each column. The (i, j) bulb is lit if $b_{ij} + r_i + c_j$ is odd.

   What is the greatest $f(n)$ such that for any setting to the b_{ij}'s, you can set the row and column switches to light at least $n^2/2 + f(n)$ bulbs?

Now, we haven't yet defined either random variables or expectations, but I think you likely already have a feel for these concepts, so let's see how *linearity of expectation* already resolves several of our appetizers. If you're not sure how to be rigorous about this, no worries, we'll proceed more methodically in the next lecture.

(1): Let the tabletop be the rectangle $[-a, a] \times [-b, b]$. Set $r = \sqrt{a^2 + b^2}$. Choose θ uniformly in $[0, \pi)$ and z uniformly in $[-r, r]$. Lay the ruler along the affine line of points (x, y) satisfying $x \cos\theta + y \sin\theta = z$. Count the number of times the ruler crosses the string.

Since we have in mind a physical string, mathematically we can model it as differentiable, and therefore the number of intersections is equal to the limit in which we decompose the string into short straight segments. A ruler can intersect such a segment only 0 or 1 times (apart from a probability 0 event of aligning perfectly). Now, our process with the ruler is such that no matter where on the table or in what orientation a straight segment lies, the probability of the ruler intersecting it is proportional to the length of the segment.

Applying linearity of expectation, we conclude that the total length is proportional to the expected number of intersections. We skip calculating the constant of proportionality.

(2): The probability that each gentleman gets his hat back is $1/N$. So the expected number of his hats that he receives (this can be only 0 or 1) is $1/N$. By linearity of expectation, the expected number of hats restored to their proper owners overall, is $\sum_1^N 1/N = 1$. (Note, if you know about independence of random events, the events corresponding to success of the various gentlemen are not independent! But that doesn't matter to since we are only adding expectations.)