Math 217 Probability and Statistics

Prof. D. Joyce, Clark University

Wednesday, 5 Sep 2006, fourth class

Due Today. From section 1.2, exercises 1–7.

Read for Friday. Section 2.1 which introduces the idea of continuous probability distributions.

Due Monday. From section 1.2, exercises 9, 11, 12, 14, 20, 25, 31a.

Today. Besides discussing homework exercises, we'll look at a couple more basic properties of probability, uniform finite probability distributions and combinatorics, a couple of examples, and note the relation between odds and probability, and perhaps introduce continuous distributions.

A couple more properties of probabilities. To say that the sample space Ω is *partitioned* into events $A_1, A_2, \ldots, A_n$ means that each outcome ω belongs to exactly one of the events $A_1, A_2, \ldots, A_n$. That's logically equivalent to saying Ω is the disjoint union of the $A_1, A_2, \ldots, A_n$. We'll have that situation from time to time, and we can use it to our advantage to compute probabilities, since we have, in that situation,

$$P(E) = P(E \cap A_1) + P(E \cap A_2) + \cdots + (E \cap A_n)$$

for any event E.

In particular, if A is any event, then A and its complement $\tilde{A}$ partition Ω, and therefore

$$P(E) = P(E \cap A) + P(E \cap \tilde{A}).$$

Uniform finite probability distributions. We've looked at a couple of these already, for instance, flipping a fair coin, or tossing a fair die. In general, a uniform distribution on a finite sample space Ω with n outcomes assigns to each outcome ω the same value $m(\omega) = \frac{1}{n}7$. Therefore, the probability of an event E is the number of outcomes in E divided by n:

$$P(E) = \frac{|E|}{|\Omega|} = \frac{|E|}{n}.$$

That means for uniform finite probabilities, you can figure out the probabilities if you can count the outcomes in an event.

The field of combinatorics is the study of counting finite sets. Central to the field is are the concepts of permutations and combinations. You've probably already seen them before, and we'll devote all of chapter 3 to combinatorics. But let's look at one exercise right now, to get a preview of what's coming up. We'll discuss exercise 7 from section 3.1: five people get on an elevator that stops at five floors. Assuming that each has an equal probability of going to any one floor, find the probability that they all get off at different floors.

Odds. Back at the time of Pascal and Fermat, two of the early researchers of probability in the 17th century, probabilities like we've been discussing with values between 0 and 1 were not used. Rather, probabilities were expressed in terms of 'odds.' These odds were usually not abstract numbers, but monetary payoffs for bets.

Let's take a concrete example. Suppose we're tossing a pair of dice repeatedly until either a sum of 7 comes up or a sum of 5 comes up. On a single toss of a pair of dice, a sum of 7 occur up 6 different ways $(1+6, 2+5, 3+4, 4+3, 5+2,$ or $6+1)$, while a sum of 5 can occur in only 4 different ways $(1+4, 3+2, 2+3,$ or $4+1)$. In this game, there are 10 outcomes, which, assuming the dice are fair, all have the same probability. Since 6 are in the event $E = \{\text{sum is } 7\}$, while 4 are in its complement $\tilde{E} = \{\text{sum is } 4\}$, we can conclude $P(E) = \frac{6}{10} = \frac{3}{5}$, while $P(\tilde{E}) = \frac{4}{10} = \frac{2}{5}$.

We can express this result in terms of odds. Since E includes 6 outcomes, and $\tilde{E}$ includes 4 outcomes, and each outcome is equally likely, we can say the odds of E to $\tilde{E}$ are in the ratio 6 to 4, or, in lowest terms, 3 to 2.

Suppose you bet on E and your opponent bets on $\tilde{E}$. Then a fair bet is where you bet \$3, and your opponent bets \$2. If the game results in E, you collect all \$5, but if it results in $\tilde{E}$, your opponent collects all \$5. Thus, for you, \$3 will get you \$2, and for your opponent, \$2 will get your opponent \$3.

It's probably much easier to deal with straight probabilities than with odds for two reasons. One is that probabilities are single numbers between 0 and 1 whereas odds are ratios of two numbers. The other is that probabilities don't involve money, wagers, and payments, but odds do, at least the historical origins of odds involve money, wagers, and payments.

The problem with continuous probability. So far, we've looked at discrete probability distributions. For a discrete distribution, each outcome has a particular positive probability, and the sum of all these probabilities is 1. When we look at continuous probability distributions, each outcome will have a probability of 0, and there's no way when you add 0s together, even an infinite number of them, you can get anything other than 0. That's the problem.

An example. Our first example is a uniform continuous distribution on the unit interval $[0,1)$. (Whether or not 1 is included is irrelevant, so don't worry whether it's a half-open interval $[0,1)$ that doesn't include 1 or a closed interval $[0,1]$ that does include 1.)

For an experiment that gives this distribution, consider spinning an arrow about the origin so that the angle that it stops at is any angle between 0° and 360°. Normalize the answer by dividing by 360° so we can treat the angle as a number X between 0 and 1. When a random variable like this X only has real numbers as values, it's called a *real random variable.*

Now, there are infinitely many possible angles that the spinner can land on. In other words, once it's normalized, the experiment gives any real number X between 0 and 1. (Strictly speaking, we should exclude 1 since that corresponds to 360°, the same angle as 0°.) And, we want any real number in the interval $[0,1)$ to have the same probability. That means that the probability that X equals any particular number has to be 0. So, how do we treat this situation?

We can still assign positive probabilities to many events. For instance, we want $P(X \leq \frac{1}{2})$ to be $\frac{1}{2}$, and we want $P(X \geq \frac{1}{2})$ to be $\frac{1}{2}$, too. In fact, if $[a,b)$ is any subinterval of $[0,1)$ we want $P(a \leq X < b)$ to be $b-a$, the length of the interval.

It turns out, as we'll soon see, that an entire continuous distribution is determined by the values of $F(x) = P(X < x)$ for all numbers x. Note how we're using the symbol X for a real random variable, that is, the numerical outcome of an experiment, but we're using the symbol x to denote particular real numbers. The function F is called a *cumulative distribution function* or, more simply, a *distribution function.*