Lecture 7 - Expected Value

6.042 - February 27, 2003

This lecture concerns the *expected value* of a random variable. Informally, this is the average value that a random variable takes on when an experiment is repeated many times. Knowing an expected value can be really handy, because this one number conveys a lot of information about a random variable. For example, suppose that we take the quiz 1 score of a 6.042 student selected uniformly at random. The expected value of that random variable is just the class average. If you know that one number, you have a pretty good idea how you did relative to the class as a whole, even though you don't know the distribution of scores in full detail.

1 Two Equivalent Definitions

There are two equivelent definitions of the expected value of a random variable. Sometimes one is more convenient, sometimes the other.

Definition 1 *The* expected value *of a random variable R over a sample space S is denoted $Ex[R]$ and defined by:*

$$Ex[R] = \sum_{x \in S} R(x) \cdot \Pr\{x\}$$

or, equivalently, by:

$$Ex[R] = \sum_{r \in R(S)} r \cdot \Pr\{R = r\}$$

The expression $R(S)$, used above, is shorthand for the set $\{R(x) \mid x \in S\}$; that is, the set of all the different values that the random variable R takes on over the whole sample space S.

The expected value of a random variable is also sometimes called the *expectation*, the *average*, or the *mean*.

1.1 The Expected Value of One Die

Suppose we roll a fair, six-sided die. Let the random variable R be the number that comes up. We can compute the expected value of R using either definition; in this case, they amount to about the same thing. Let's use the second one:

$$
\begin{aligned}
\operatorname{Ex}[R] &= \sum_{i=1}^{6} i \cdot \Pr\{R = i\} \\
&= 1 \cdot \frac{1}{6} + 2 \cdot \frac{1}{6} + 3 \cdot \frac{1}{6} + 4 \cdot \frac{1}{6} + 5 \cdot \frac{1}{6} + 6 \cdot \frac{1}{6} \\
&= \frac{7}{2}
\end{aligned}
$$

Thus, the expected value of the number rolled on a fair die is $\frac{7}{2}$. This means that if you roll a die many times, you should expect the average value to be about $\frac{7}{2}$. Obviously, you should not expect the value $\frac{7}{2}$ ever to come up on an individual roll!

1.2 Equivalence of the Two Definitions

Let's prove that the two definitions of expected value given at the beginning of this section actually are equivalent. In the first definition, there is one term in the summation for each outcome in the sample space. By grouping together outcomes on which the random variable R takes the same value, we get the second definition. Here are the gory details:

$$
\begin{aligned}
\sum_{x \in S} R(x) \cdot \Pr\{x\} &= \sum_{r \in R(S)} \sum_{x \in S : R(x) = r} R(x) \cdot \Pr\{x\} \\
&= \sum_{r \in R(S)} \sum_{x \in S : R(x) = r} r \cdot \Pr\{x\} \\
&= \sum_{r \in R(S)} r \cdot \left(\sum_{x \in S : R(x) = r} \Pr\{x\} \right) \\
&= \sum_{r \in R(S)} r \cdot \Pr\{R = r\}
\end{aligned}
$$

Here the notation $x \in S : R(x) = r$ refers to all outcomes x in the sample space S such that $R(x) = r$. After the first step, we're summing up exactly the same terms as before, but indexing them a different way. The second step uses the fact that $R(x) = r$ for all terms in the inner summation. In the third step, we pull r out of the inner sum. This inner sum is then the very definition of the probability of the event $R = r$, so we replace it by $\Pr\{R = r\}$ in the last step. Since one can equally well reason from the last expression back to the first, the two definitions are equivalent.

2 Chuck-a-Luck

Here's yet another fun 6.042 game! You pick a number between 1 and 6. Then you roll three fair, independent dice.

- If your number never comes up, then you lose a dollar.
- If your number comes up once, then you win a dollar.
- If your number comes up twice, then you win two dollars.
- If your number comes up three times, you win three dollars!

Is this a good game to play? That is, if you play many times, do you expect to win money overall?

We can formulate these questions mathematically by defining a random variable R, which denotes the amount of money that you win in a game of Chuck-a-Luck. Then, for example, we have $R = 1$ when your number comes up once and $R = -1$ when it never comes up. Asking whether you should expect to win money playing Chuck-a-Luck for a long time is equivalent to asking whether the expectation of R is positive; that is, is $\mathrm{Ex}[R] > 0$?

Tree diagrams can often be useful for computing expectations as well as probabilities. Figure 1 shows a tree diagram for Chuck-a-Luck. The branches marked with probability $\frac{1}{6}$ correspond to die rolls where your number comes up; those marked $\frac{5}{6}$ correspond to rolls where your number does not come up. In addition to noting the probability of each outcome, we've indicated the value of the random variable R as well. This is precisely the information we need to compute the expected value of R:

$$\begin{aligned}
\mathrm{Ex}[R] &= \sum_{r \in R(S)} r \cdot \Pr\{R = y\} \\
&= -1 \cdot \left(\frac{125}{216}\right) + 1 \cdot \left(3 \cdot \frac{25}{216}\right) + 2 \cdot \left(3 \cdot \frac{5}{216}\right) + 3 \cdot \left(\frac{1}{216}\right) \\
&= -\frac{17}{216}
\end{aligned}$$

If you play Chuck-a-Luck many times, you should expect to lose about 8 cents per game!

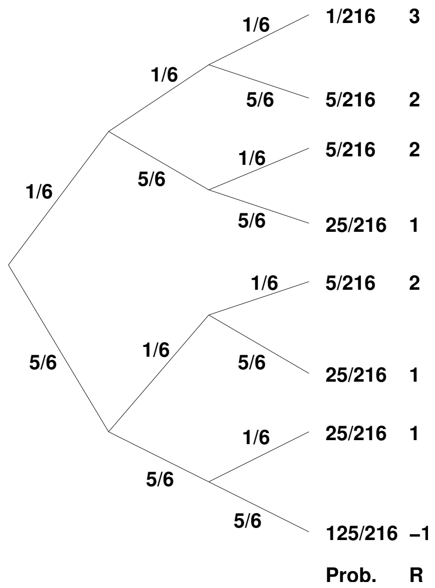

Figure 1: Tree diagram for Chuck-a-Luck.

3 Linearity of Expectation

Expectation obeys a simple, but remarkably powerful rule called *linearity of expectation*. This rule is stated and proved below.

Theorem 2 *For all random variables X and Y and all real numbers a and b:*

$$Ex[aX + bY] = aEx[X] + bEx[Y]$$

Proof. Let S be the sample space. The left side can be transformed into the right side as follows:

$$\begin{aligned}
\text{Ex}[aX + aY] &= \sum_{s \in S} (aX(s) + bY(s)) \cdot \Pr\{s\} \\
&= \sum_{s \in S} aX(s) \cdot \Pr\{s\} + \sum_{s \in S} bY(s) \cdot \Pr\{s\} \\
&= a \sum_{s \in S} X(s) \cdot \Pr\{s\} + b \sum_{s \in S} Y(s) \cdot \Pr\{s\} \\
&= a\text{Ex}[X] + b\text{Ex}[Y]
\end{aligned}$$

The first step uses the definition of expectation. In the second step, we split the single sum into two. Then we pull the constants a and b out of their summations and finish by applying the definition of expectation again. □

Linearity of expectation is most often used to compute the expected value of a sum of random variables. In that case, the constants a and b are both equal to 1, and the theorem says that the expectation of a sum of random variables is the sum of their expectations:

$$\mathrm{Ex}[X+Y] = \mathrm{Ex}[X] + \mathrm{Ex}[Y]$$

The great thing about linearity of expectation is that *no independence is required.* This is really useful, because dealing with independence is a pain, and we often need to work with random variables that are not independent.

3.1 Expected Value of Two Dice

What is the expected value of the sum of two fair dice?

Let the random variable X be the number on the first die, and let Y be the number on the second. We observed earlier that the expected value of one die is $\frac{7}{2}$. We can find the expected value of the sum using linearity of expectation:

$$\begin{aligned}
\mathrm{Ex}[X+Y] &= \mathrm{Ex}[X] + \mathrm{Ex}[Y] \\
&= \frac{7}{2} + \frac{7}{2} \\
&= 7
\end{aligned}$$

Notice that we did *not* have to assume that the two dice were independent. The expected sum of two dice is 7, even if they are glued together! (This is provided that gluing does not change weights to make the individual dice unfair.)

Proving that the expected sum is 7 with a tree diagram would be hard; there are 36 cases. And if we did not assume that the dice were independent, the job would be a nightmare!

3.2 The Hat-Check Problem

There is a dinner party where n men check their hats. The hats are mixed up during dinner, so that afterward each man receives a random hat. In particular, each man gets his own hat with probability $1/n$. What is the expected number of men who get their own hat?

Without linearity of expectation, this would be a very difficult question to answer. We might try the following. Let the random variable R be the number of men that get their own hat. We want to compute $\mathrm{Ex}[R]$. By the definition of expectation, we have:

$$\operatorname{Ex}[R] = \sum_{k=0}^{n} k \cdot \Pr\{R = k\}$$

Now we are in trouble, because evaluating $\Pr\{R = k\}$ is a mess and we then need to substitute this mess into a summation. Furthermore, to have any hope, we would need to fix the probability of each permutation of the hats. For example, we might assume that all permutations of hats are equally likely, even though this is not in the statement of the problem. As a result, our solution would apply only in this restricted circumstance.

So let's use linearity of expectation. As before, let the random variable R be the number of men that get their own hat. The first trick is to express R as a sum of indicator variables. In particular, let R_i be an indicator for the event that the i-th man gets his own hat. That is, $R_i = 1$ is the event that he gets his own hat, and $R_i = 0$ is the event that he gets a wrong hat. The number of men that get their own hat is the sum of these indicators:

$$R = R_1 + R_2 + \cdots + R_n$$

These indicator variables are *not* mutually independent. For example, if $n-1$ men all get their own hats, then the last man is certain to receive his own hat. So R_n is not independent of the other indicator variables. But, since we plan to use linearity of expectation, we *don't care* whether the indicator variables are independent. No matter what, we can take the expected value of both sides of the equation above and apply linearity of expectation:

$$\begin{aligned}
\operatorname{Ex}[R] &= \operatorname{Ex}[R_1 + R_2 + \ldots + R_n] \\
&= \operatorname{Ex}[R_1] + \operatorname{Ex}[R_2] + \ldots + \operatorname{Ex}[R_n]
\end{aligned}$$

Now our problem is to compute the expected value of each indicator variable R_i. To do that, let's go back to the definition of expectation, bearing in mind that an indicator variable takes on only the values 0 and 1.

$$\begin{aligned}
\operatorname{Ex}[R_i] &= \sum_{r \in R_i(S)} r \cdot \Pr\{R_i = r\} \\
&= 0 \cdot \Pr\{R_i = 0\} + 1 \cdot \Pr\{R_i = 1\} \\
&= \Pr\{R_i = 1\}
\end{aligned}$$

The calculation above illustrates a general rule: *the expectation of an indicator is equal to the probability that the indicator is 1.*

In this case, the indicator R_i is equal 1 precisely when the i-th man gets his hat back. According to the problem statement, this happens with probability $1/n$. Therefore, we have:

$$\begin{aligned}
\mathrm{Ex}[R] &= \mathrm{Ex}[R_1] + \mathrm{Ex}[R_2] + \ldots + \mathrm{Ex}[R_n] \\
&= \Pr\{R_1 = 1\} + \Pr\{R_2 = 1\} + \ldots + \Pr\{R_n = 1\} \\
&= \frac{1}{n} + \frac{1}{n} + \ldots + \frac{1}{n} \\
&= 1
\end{aligned}$$

On average, exactly one man gets his own hat back. Amazingly, this is true whether only two men went to dinner or a thousand!

4 A Bar Game

Three guys in a bar are playing a game. First, each player puts two dollars into the pot. Then each player secretly writes "heads" or "tails" on a napkin. The bartender then flips a fair coin. The players reveal their napkins and the pot is divided evenly among all players that correctly predicted the outcome of the coin toss. If no one predicted correctly, everyone takes their money back.

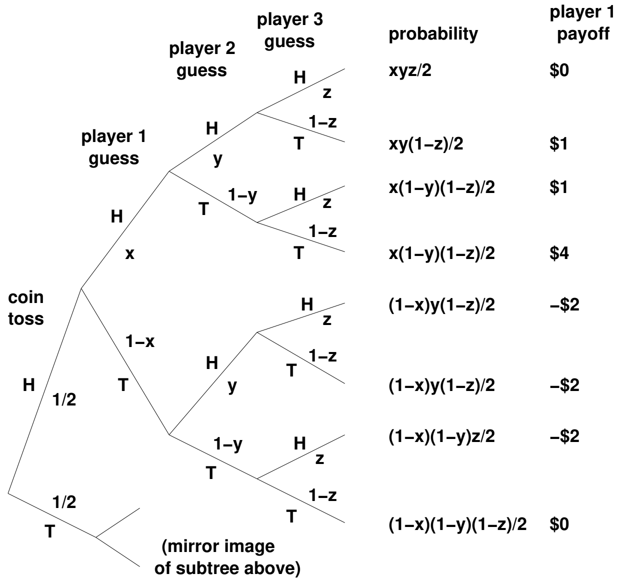

Figure 2: Tree diagram for the bar game.

On first inspection, this appears to be a perfectly fair game. Since the roles of the three players are symmetric, there is no way for one player to get a special advantage.

But let's look more closely. Let x, y, and z be the probabilities that the first, second, and third players choose "heads". The expected payoff for the first player is computed via the tree diagram in Figure 2 and some routine, but messy algebra. The conclusion is:

$$\text{Ex}\,[\text{player 1 payoff}] \quad = \quad x(1-y-z) + 2yz - \frac{y+z}{2}$$

This equation requires some interpretation. Initially, suppose that the second and third players are both equally likely to predict heads and tails ($y = z = \frac{1}{2}$). Then the expression above is zero, meaning that the first player breaks even.

But now suppose that the second player always guesses heads ($y = 1$) and the third player always guesses tails ($z = 0$) or vice versa. Either way, we have $y + z = 1$ and $yz = 0$. In this case, the expected payoff for the first player is $-\frac{1}{2}$. In other words, if the second and third players secretly collude by always betting oppositely, the first player will lose an average of fifty cents per game, no matter what he does!