Lecture 22 - Random Processes II

6.042 - May 8, 2003

1 Coupon Collecting

Each box of cereal contains a coupon. There are n different kinds of coupon, and all are equally likely to appear in each box of cereal. If you collect one coupon of each kind, then you can mail them all in to the manufacturer and receive a prize in exchange. How many boxes of cereal must you eat to earn the prize? This is a classic puzzle known as the *coupon collector problem*.

First, some preliminaries. A probabilistic experiment that either succeeds or fails is called a *Bernoulli trial*. We need a fact that, while stated in terms of Bernoulli trials, was actually proved in the previous lecture:

Fact 1 *Suppose that we perform a sequence of independent Bernoulli trials, each of which succeeds with probability p. Then the expected number of trials needed to obtain one success is $1/p$.*

Let's first think about the coupon-collecting process informally. After eating one box of cereal, you have one kind of coupon. The second box of cereal probably contains a different coupon, giving you two kinds. But the second box *might* contain the same kind of coupon as the first box, leaving you still with only one kind. But, over time, you eventually acquire a second kind of coupon, and then a third, and then a fourth, and so on. However, acquiring new kinds of coupon becomes progressively harder; increasingly often, you open a box of cereal to discover a coupon that you already have. Eventually, however, you get them all.

Now let's recast this informal description into a mathematical analysis and solve the problem. The key is to define just the right set of random variables. In particular, let the random variable X_k be *the number of additional boxes of cereal that you must eat to get a new kind of coupon, when you already have $k-1$ different kinds of coupon.* Then define:

$$X = X_1 + X_2 + \ldots + X_n$$

Thus, X is the number of boxes you must eat to get one kind of coupon, plus the number of additional boxes you must eat to get a second kind of coupon, plus the number of additional boxes to get a third kind, and so forth. In other words, X is the total number of boxes of cereal that you must eat in order to win the prize. Therefore, our goal is to compute $\text{Ex}[X]$. By linearity of expectation, we have:

$$
\begin{aligned}
\mathrm{Ex}[X] &= \mathrm{Ex}[X_1 + X_2 + X_3 + \ldots + X_n] \\
&= \mathrm{Ex}[X_1] + \mathrm{Ex}[X_2] + \mathrm{Ex}[X_3] + \ldots + \mathrm{Ex}[X_n]
\end{aligned}
$$

All that remains is to compute $\mathrm{Ex}[X_k]$, where $1 \leq k \leq n$. So suppose that you have $k-1$ coupons. We can now regard opening a cereal box as a Bernoulli trial where "success" means getting a new kind of coupon. With probability $(k-1)/n$, you fail; that is, you get a kind of coupon that you already possess. Therefore, you succeed with probability $(n-k+1)/n$. Thus, by Fact 1, the expected number of boxes that you must eat to get a new kind of coupon is $n/(n-k+1)$. Plugging this observation into the equation above gives:

$$
\begin{aligned}
\mathrm{Ex}[X] &= \mathrm{Ex}[X_1] + \mathrm{Ex}[X_2] + \mathrm{Ex}[X_3] + \ldots + \mathrm{Ex}[X_n] \\
&= \frac{n}{n} + \frac{n}{n-1} + \frac{n}{n-2} + \ldots + \frac{n}{1} \\
&= n \cdot \left(\frac{1}{n} + \frac{1}{n-1} + \frac{1}{n-2} + \ldots + \frac{1}{1} \right) \\
&= nH_n
\end{aligned}
$$

In the third step, we pull an n out of each term. For the last step, note that the sum in parentheses is equal to the n-th harmonic number, H_n. Usually, the terms in a harmonic sum are listed in the reverse order, but addition is commutative, of course.

Therefore, you must eat an average of $nH_n \approx n \ln n$ boxes of cereal to collect all the coupons and win the prize. This solution has remarkably many applications. For example, suppose you want to meet at least one person born on each day of the year. How many people must you meet? Ignoring leap days, as usual, the expected number is:

$$
365 \cdot H_{365} \approx 2365
$$

Even more problems can be solved by variations on our general approach to the coupon collector problem: break a random process into a sequence of analyzable steps and then use linearity of expectation to assemble the results into an analysis of the overall process.

2 The Doubling Scheme

Suppose that you and I bet \$5 on the outcome of a fair coin toss. If you win, you walk away \$5 richer. But if you lose, we bet \$10 on the outcome of a second coin toss. If you win the second bet, you walk away $10 - 5 = 5$ dollars richer. But if you lose the second bet, we bet \$20 on a third coin toss. If you win the third bet, you walk away $20 - 10 - 5 = 5$ dollars richer. But if you lose, we keep going, doubling the stakes with each bet until you eventually win. Your earnings from that single win more than offset all your preceding losses, and you walk away \$5 richer. No matter what, you make 5 bucks!

Every day, this doubling scheme is reinvented by a roulette player somewhere in the world. But, as you may have noticed, the world's casinos aren't broke. As a practical matter, the problem with this scheme is simple: most of the time, you make \$5. However, eventually you have a streak of bad luck and either go broke or hit the house's upper betting limit. Either way, you have lost a big pile of money, which you can not recover by doubling your bet again.

2.1 A Faulty Expectation Analysis

Even though the doubling scheme doesn't work in practice, a theoretical analysis is a stress-test for our probabilistic machinery. Let's try to compute your expected profit from the doubling scheme. Number the bets 0, 1, 2, 3, ... Let the random variable X_k be your payoff from the k-th bet. Then we have:

$$X_k = \begin{cases} 0 & \text{with probability } 1 - \frac{1}{2^k} \text{ (you've already gone home)} \\ 5 \cdot 2^k & \text{with probability } \frac{1}{2} \cdot \frac{1}{2^k} \text{ (you win)} \\ -5 \cdot 2^k & \text{with probability } \frac{1}{2} \cdot \frac{1}{2^k} \text{ (you lose)} \end{cases}$$

The first case above corresponds to the possibility that you won an earlier bet and already went home; thus, you have nothing to win or lose from the k-th bet. The second and third cases correspond to the possibilities of winning and losing the k-th bet. Note that the expected value of X_k is zero; this makes sense, since this is your expected payoff from an even bet on a fair coin, if you bet at all.

Your overall profit from the doubling scheme is the sum of your payoffs for every bet. Thus, we can compute your expected overall profit using linearity of expectation:

$$\begin{aligned} \operatorname{Ex}\left[\sum_{k=0}^{\infty} X_k\right] &= \sum_{k=0}^{\infty} \operatorname{Ex}\left[X_k\right] \\ &= 0 \end{aligned}$$

This is bad news! We argued earlier that you were certain to win \$5 eventually. But then how can your expected payoff be zero?

2.2 Expectation of an Infinite Sum

Long ago, in the before-time, we proved that the the expectation of the sum of two random variables is equal to the sum of their expectations:

$$\mathrm{Ex}[A+B] \quad = \quad \mathrm{Ex}[A] + \mathrm{Ex}[B]$$

By induction, we can prove that the sum of any finite number of random variables is equal to the sum of their expectations:

$$\mathrm{Ex}[A_1 + A_2 + \ldots + A_n] \quad = \quad \mathrm{Ex}[A_1] + \mathrm{Ex}[A_2] + \ldots + \mathrm{Ex}[A_n]$$

But, so far, we have no reason to believe that the expectation of an *infinite sum* of random variables is equal to the sum of their expectations. This was precisely what we relied upon in our argument in the preceding section, and this was our error.

In certain cases, linearity of expectation does hold for infinite sums.

Theorem 1 *Let $X_0, X_1, X_2, \ldots$ be a (possibly infinite) sequence of random variables such that*

$$\sum_{k=0}^{\infty} Ex[|X_k|]$$

converges. Then:

$$Ex\left[\sum_{k=0}^{\infty} X_k\right] \quad = \quad \sum_{k=0}^{\infty} Ex[X_k]$$

Proof. We can reason as follows:

$$\begin{aligned}
\sum_{k=0}^{\infty} \mathrm{Ex}\left[X_k\right] &= \sum_{k=0}^{\infty} \sum_{s \in S} X_k(s) \cdot \Pr\{s\} \\
&= \sum_{s \in S} \sum_{k=0}^{\infty} X_k(s) \cdot \Pr\{s\} \\
&= \sum_{s \in S} \left(\Pr\{s\} \cdot \sum_{k=0}^{\infty} X_k(s) \right) \\
&= \mathrm{Ex}\left[\sum_{k=0}^{\infty} X_k \right]
\end{aligned}$$

The first step uses the definition of expectation. In the second, we swap the sums, which is valid because of our absolute convergence assumption. Next, we pull $\Pr\{s\}$ out of the inner summation. The last step uses the definition of expectation once more. $\square$

Note that the absolute convergence condition in the theorem statement does not hold for the doubling scheme. There, we have:

$$\begin{aligned}
\sum_{k=0}^{\infty} \mathrm{Ex}\left[|X_k|\right] &= \sum_{k=0}^{\infty} \frac{1}{2^k} \cdot (5 \cdot 2^k) \\
&= \sum_{k=0}^{\infty} 5
\end{aligned}$$

The last sum is clearly not convergent. Thus, the analysis of the doubling scheme using linear of expectation was wrong. Your expected payoff from the doubling scheme really is \$5. But remember that this only works if both you and the person you're betting against have an unlimited amount of money to bet. And, if that's the case, why try to make \$5?

3 The Truel

Three gunfighters meet for a *truel*, a three-person duel. Gunfighter A hits his target 50% of the time, gunfighter B hits 75% of the time, and gunfighter C hits 100% of the time. The gunfighters take turns shooting in the order A, B, C, A, B, C, etc. Of course, a dead gunfighter misses his turn. The last one standing is the winner.

What is A's best strategy? If A kills C, then B will probably kill A on the next shot. On the other hand, if A kills B, then C will certainly kill A on the next shot. This does not look good. But there is a another possibility: A could intentionally miss, let B and C shoot it out, and then try to kill the winner! Let's evaluate that strategy, assuming that B and C actually try to hit each other.

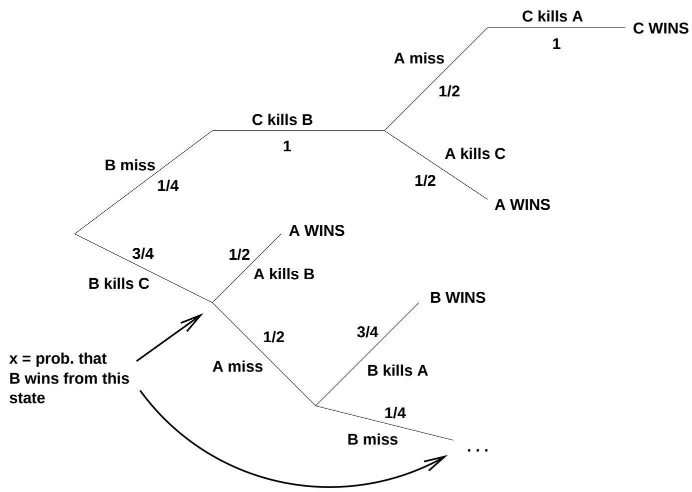

From the tree diagram, we have:

$$\begin{aligned}\Pr\{\text{C wins}\} &= \frac{1}{4}\cdot 1\cdot\frac{1}{2}\cdot 1\\ &= \frac{1}{8}\\ &= 12.5\%\end{aligned}$$

Now let x be the probability that B eventually wins in the situation where C is dead and A has the next shot. This situation arises at two different points in our tree diagram. We can exploit that fact to obtain an equation expressing x in terms of itself:

$$x = \frac{1}{2}\cdot\frac{3}{4}+\frac{1}{2}\cdot\frac{1}{4}\cdot x$$

Solving this equation, we find that $x = 3/7$. The probability that B wins overall is:

$$\begin{aligned}\Pr\{\text{B wins}\} &= \frac{3}{4}\cdot x\\ &= \frac{9}{28}\\ &\approx 32.1\%\end{aligned}$$

Finally, we have:

$$\begin{aligned}\Pr\{\text{A wins}\} &= 1 - \Pr\{\text{B wins}\} - \Pr\{\text{C wins}\} \\ &\approx 55.4\%\end{aligned}$$

Amazingly, the worst shooter has the best chance of winning, and the best shooter has the worst chance of winning!

Of course, an explicit assumption in this analysis was that B and C are both shooting to kill, unlike A in the first round. If B and C have no such requirement, then the problem is underspecified; there is no definite mathematical solution. Every gunfighter might reason that he is better off not shooting and the whole lot might go toast smores over a campfire.