Intermediate Math Circles
March 28, 2012
Probability: Expectation and Ruination

Last week we talked about conditional probabilities, like "if we pick a two-child family at random, if one of the children is a girl, then the conditional probability that both are girls is 1/3." We also saw the *normal distribution*: when you repeatedly add up numbers sampled from any distribution, the distribution of the sum converges to a normal distribution, described by its expected value (maximum location/horizontal offset) and standard deviation (stretching factor). Coins can be *fair* (50% heads, 50% tails) or *biased* (some probability $p \neq \frac{1}{2}$ of heads, $1-p$ of tails). Dice, coins, and biased coins all converge to the normal distribution, with different expected values and standard deviations.

This week we'll get to the **Gambler's Ruin** game: you repeatedly flip a fair coin; when it is heads you win \$1, and when it is tails you lose \$1. If you started the night with \$30, and you keep playing until your total money reaches either \$100 (success) or \$0 (ruin), *what is the probability of success or ruin*?

Defining Expected Value

A *random variable* is some variable that takes on different values with different probabilities. For example, the number on top of a die is a random variable T defined by

$$\mathsf{Pr}[T=1] = \frac{1}{6}, \mathsf{Pr}[T=2] = \frac{1}{6}, \mathsf{Pr}[T=3] = \mathsf{Pr}[T=4] = \mathsf{Pr}[T=5] = \mathsf{Pr}[T=6] = \frac{1}{6};$$
$$\mathsf{Pr}[T=i] = 0 \text{ for all other } i.$$

The expected value of a random variable is its average value. We'll start with the die example first. We calculate the expected value by multiplying each possible value by its probability, and adding these numbers up.

The expected value of T is

$$1 \times \frac{1}{6} + 2 \times \frac{1}{6} + 3 \times \frac{1}{6} + 4 \times \frac{1}{6} + 5 \times \frac{1}{6} + 6 \times \frac{1}{6} = \frac{21}{6} = \frac{7}{2} = 3.5.$$

The symbol for the expected value of T is $\mathsf{E}[T]$. If you wanted to write out the formula for expected value, the simplest method uses sigma sum notation,

$$\mathsf{E}[X] = \sum_v v \times \mathsf{Pr}[X=v]$$

where you add up the sum over all possible values v that X can take on.

When you flip a coin, let H be 1 if the top is "heads," and 0 if the top is "tails." What is $\mathsf{E}[H]$? The probability that H is 1 is $\frac{1}{2}$, and the probability that H is 0 is $\frac{1}{2}$. So the expected value is

$$1 \times \frac{1}{2} + 0 \times \frac{1}{2} = \frac{1}{2}.$$

Let's roll another die, but this one has the English words "one," "two," "three," "four," "five," "six" written on the sides. What is the expected number of letters on top, when you roll the die?

There are multiple ways to solve this; here's one solution. Let L be the length of the word on top. Since

$$\mathsf{Pr}[L=3]=\frac{3}{6}, \mathsf{Pr}[L=4]=\frac{2}{6}, \text{ and } \mathsf{Pr}[L=5]=\frac{1}{6},$$

we get

$$\mathsf{E}[L]=3\times\frac{3}{6}+4\times\frac{2}{6}+5\times\frac{1}{6}=22/6=3.666.$$

Let's flip 2 coins. What is the expected number of total heads on top?

Let N be this number. We have

$$\mathsf{Pr}[N=0]=\frac{1}{4}, \mathsf{Pr}[N=1]=\frac{2}{4}, \mathsf{Pr}[N=2]=\frac{1}{4}.$$

So the expected value is

$$E[N]=0\times\frac{1}{4}+1\times\frac{2}{4}+2\times\frac{1}{4}=4/4=1.$$

Finally, let's flip 3 coins. What is the expected number of heads on top?

Let N be this number. We have

$$\mathsf{Pr}[N=0]=\frac{1}{8}, \mathsf{Pr}[N=1]=\frac{3}{8}, \mathsf{Pr}[N=2]=\frac{3}{8}, \mathsf{Pr}[N=3]=\frac{1}{8},$$

so the expected value is

$$\mathsf{E}[N]=0\times\frac{1}{8}+1\times\frac{3}{8}+2\times\frac{3}{8}+3\times\frac{1}{8}=\frac{12}{8}=\frac{3}{2}.$$

Look at the expected number of heads when you flip 1, 2, or 3 coins. Does it look like there is a pattern? If you want, check whether a similar pattern holds for rolling dice or flipping biased coins.

The expected number of heads for 1 coin is $\frac{1}{2}$, for 2 coins is $\frac{2}{2}$, and for 3 coins is $\frac{3}{2}$. It looks like for k coins, the expected number of heads is $k/2$. For dice, the expected sum of 1 die is $7/2$ and of two dice is $7=2\times\frac{7}{2}$, a similar pattern.

The pattern does continue, and something much more general is true. Let A and B be two random variables, and let C be a third random variable equal to the sum of A and B. Then:

Linearity of expectation: The expected value of C equals the expected value of A, plus the expected value of B. In short, $\mathsf{E}[A+B]=\mathsf{E}[A]+\mathsf{E}[B]$.

So far we only saw examples of this where A and B are independent, but this works even if A and B are dependent. (For example, pick a random word from the dictionary, let A be the number of consonants it contains, B the number of vowels, and C the number of letters overall.)

Linearity of Expectation tells us something about repeating the same experiment N times: if you add up N independent samples of the same random variable, the expected sum equals N times the expected value of a single trial. This tells us where the limiting "peak" will show up, in normal distributions. For example, the peak when we rolled 100 dice was at the value 350, which is 100 times the expected value of a single roll.

Gambler's Ruin, and Random Walks

You are playing the following game at a casino: you flip a fair coin, and when it is heads you win \$1, and when it is tails you lose \$1. If you started the night with \$30, and you keep playing until your total money reaches either \$100 or \$0 (gambler's ruin), what is the probability of getting \$100?

Looking at small examples is a nice way to approach the problem, but it is very tricky! Even if we started with \$1 and just wanted to know about winning \$3 before being ruined, it is not easy to compute the final answer. Clearly after one step you are ruined with probability $\frac{1}{2}$ and have \$2 otherwise. If you take a second step, you get to \$3 with probability $\frac{1}{4}$, and are back to \$1 with probability $\frac{1}{4}$. Continuing, we find the probability of getting to \$3 overall is

$$\frac{1}{4} + \frac{1}{16} + \frac{1}{64} + \cdots = \frac{1}{3}.$$

But this approach is too tricky to generalize. Instead, let's look at the expected values. Note *your expected value that after the first flip* is

$$\frac{1}{2} \times 29 + \frac{1}{2} \times 31 =$$

After two flips, the expected value is

$$\frac{1}{4} \times 28 + \frac{2}{4} \times 30 + \frac{1}{4} \times 32 =$$

This continues for the rest of the game:

Lemma. When you flip a coin, the expected value of your money doesn't change.

Proof: If you have $\$N$ and flip a coin, you will win \$1 or lose \$1 each with probability 1/2. So the *expected* value after the flip is

$$\frac{1}{2}(N-1) + \frac{1}{2}(N+1) = \frac{N-1+N+1}{2} = \frac{2N}{2} = N. \qquad \square$$

This is very useful! If you do a flip every minute until you get to \$0 or \$100, this means that for all times T, the expected amount of money you have at time T is \$30.

When the game is over, we either have \$0 or \$100. Let p be the probability of success (ending with \$100). The probability of ruin must be $1 - p$. Then by the definition of expected value,

$$\$30 = \mathsf{E}[\text{money at end of game}] = \$100 \times p + \$0 \times (1-p) = \$100 \times p.$$

Finally, we solve this to get that the probability of success is $p = \$30/\$100 = 0.3$.

Random Walks

Imagine a small bead attached to the x-axis, which starts at $(30, 0)$. Each minute, it either moves rightwards on the x-axis increasing by $(+1, 0)$, or leftwards moving by $(-1, 0)$, with equal probability. If you put sticky tape at $(0, 0)$ and $(100, 0)$ to stop the bead once it gets to either point, this random process exactly models the amount of money in our Gambler's Ruin game.

What happens when the particle wanders in 2D or higher? Using an excel file from `http://excelunusual.com/a-2d-random-walk-model-the-drunk-man-animation/`, I'll show a video of 10000 random walkers in 2D: `http://www.youtube.com/watch?v=NBdBFQbsfmc`

You can see in the video that the typical distance from the origin increases with time, but the rate of increase gets smaller with time. There is a large bang at the beginning and the particles expand rapidly, but then the expansion slows down.

If you start to analyze the situation, you'll find that the typical distance from the origin increases according to the square root of time elapsed. For example, the picture after 400 steps looks a lot like a $2\times$ version of the picture after 100 steps. It's not a coincidence that $2 = \sqrt{400/100}$. This $\sqrt{N}$ phenomenon happens in 1D, 2D, 3D, and in any number of dimensions.

What does this have to do with drinking? A drunk man walking in random directions is a good model for a particle taking a random walk. The mathematician Shizuo Kakutani wisely said, "a drunk man will return home while a drunk bird might lose its way forever." This expresses a theorem of Georg Pólya: a 2D random walker eventually hits to the origin again with probability 1, but a 3D random walker has a probability of only 34% of getting back to the origin.

The $\sqrt{\text{Root}}$ of The Problem: Variance and Standard Deviation

How do you determine the other scaling parameter of normal distributions? The *variance* of a random variable is the expected value of the square of the difference from the expected value:

$$\mathsf{Var}[X] = \mathsf{E}[(X - \mathsf{E}[X])^2] \text{ which turns out to be the same as } \mathsf{Var}[X] = \mathsf{E}[X^2] - \mathsf{E}[X]^2.$$

The *standard deviation* is the square root of the variance. It is not hard to prove, but we won't do it here, that the standard deviation of N independent copies of a random variable equals $\sqrt{N}$ times the original. For example, in rolling a die and looking at the number T on top, we get that the variance is

$$\mathsf{Var}[T] = \mathsf{E}[(T - \mathsf{E}[T])^2] = \frac{1}{6}(1-\tfrac{7}{2})^2 + \frac{1}{6}(2-\tfrac{7}{2})^2 + \frac{1}{6}(3-\tfrac{7}{2})^2 + \frac{1}{6}(4-\tfrac{7}{2})^2 + \frac{1}{6}(5-\tfrac{7}{2})^2 + \frac{1}{6}(6-\tfrac{7}{2})^2$$

$$= (25/4 + 9/4 + 1/4 + 1/4 + 9/4 + 25/4)/6 = 70/24$$

so the standard deviation of one roll is $\sqrt{70/24} \approx 1.71$ and the standard deviation of 100 rolls is $\sqrt{100} \cdot \sqrt{70/24} \approx 17.1$. These predictions — a mean of 350 and a standard deviation of 17.1 — agree with the maximum and inflection points on the diagram from last week.

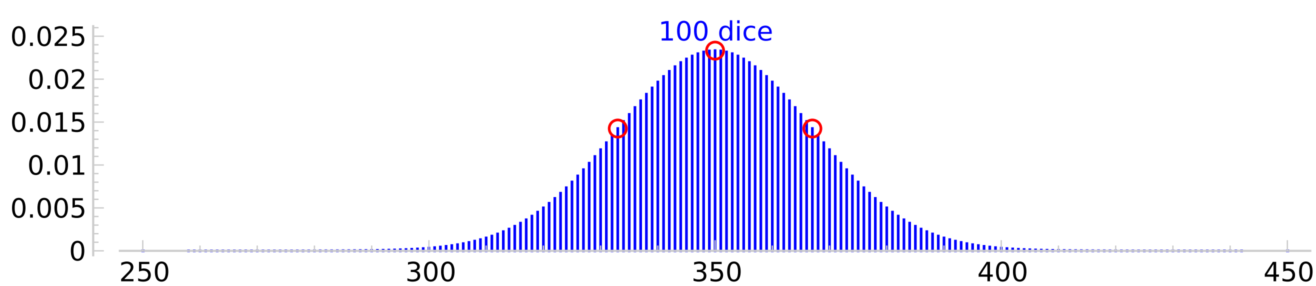

The $\sqrt{N}$ factor has many important consequences: for example, if you want to re-run an experiment and reduce the error by a factor of 10, you need to collect 100 times as many data points.

Problems

1. **Simpson's Paradox.** In January, Homer won 1 out of 5 games of StarCraft. In February, he won 3 out of 4 games. His wife Marge won 2 out of 9 games in January and 4 out of 5 games in February. Who was the better player each month? Who was better overall?

2. You are trapped on an desert island, gambling with your friend. You only have one coin, and it is biased. You don't even know what probability it has to come up heads. How can you use this biased coin to simulate a fair coin?

3. **The St. Petersburg Paradox.** You are playing a new game at the casino. You keep flipping a fair coin repeatedly, until it turns up heads. If you got k tails before reaching the first heads, then you win 2^k dollars. What is a fair price for the casino to charge for this game — what is the expected value that you win?

4. You are playing rock-paper-scissors, getting \$1 when you win, paying \$1 when you lose, and doing nothing when you tie. A *strategy* is a random variable whose only possible values are rock, paper, scissors. A strategy could also assign probability 0 to some actions.

   - Find a strategy to use so that no matter what your opponent picks, your expected value from each round is non-negative.
   - Show there is only one such strategy.
   - You are allergic to paper and must assign it probability 0. What strategy minimizes the expected amount of money you will lose, no matter what your opponent picks?

5. In the *crown and anchor* game, there are 6 symbols (the 4 card suits, crown, and anchor). When you bet \$1 on a symbol, you roll three dice that have the 6 symbols on their 6 sides. If your symbol appears on top of $T \geq 1$ dice you win $(T+1)$ dollars, otherwise you lose your \$1 investment. Show that your expected profit each round is $-\$17/216$.

6. Prove $\mathsf{E}[(X - \mathsf{E}[X])^2] = \mathsf{E}[X^2] - \mathsf{E}[X]^2$. Prove that this implies $\mathsf{E}[X^2] \geq \mathsf{E}[X]^2$, which is a version of the *arithmetic mean-quadratic mean* inequality.

7. Consider the sequences (H,T,T) and (H,T,H). If you keep flipping a coin, which sequence has a smaller expected number of flips before it shows up? Watch the talk:
   `http://www.ted.com/talks/peter_donnelly_shows_how_stats_fool_juries.html`

8. Consider *Gambler's ruin with an unfair coin*: on every flip you win \$1 with probability $p \neq \frac{1}{2}$, and lose \$1 with probability $1-p$. You start with \$30, and stop once you hit \$0 or \$100. **What is the probability of success?**
   To compute this, if you have $\$x$ in real life, define your *virtual value* to be $((1-p)/p)^x$. Show that flipping a coin doesn't change your expected virtual value. Then, using the same ideas as we used for normal Gambler's ruin, show that the probability of success is

$$\frac{[(1-p)/p]^{30} - 1}{[(1-p)/p]^{100} - 1}.$$

Answers to Problems

1. In January, Homer won 20% of games and Marge won 22.2%, so Marge was better. In February, Homer won 75% of games and Marge won 80%, so Marge was better. But overall, Homer won 44.4% and Marge only won 42.8%. Marge won both months, but Homer won overall! This situation does arise in statistics, and the right way to resolve the apparent contradiction depends on the application. See "Simpson's Paradox" on Wikipedia, and the article "Simpson's paradox: An Anatomy" by Judea Pearl.

2. Let p be the probability that this unfair coin lands heads-up. If you flip the coin twice, notice that there are four possible outcomes: HH, HT, TH, TT where XY means X on the first flip, and Y on the second flip. The important thing is that HT and TH both have equal probability $p(1-p) = (1-p)p$. Therefore, flip the coin twice, treat HT like a "fair head" and TH like a "fair tail." If you got HH or TT, start over. In this way, you eventually will get either a fair head or a fair tail, each with equal probability.

3. Notice that the probability of winning 2^k dollars is $(\frac{1}{2})^{k+1}$. So, the expected value is

$$\frac{1}{2} \times 1 + \frac{1}{4} \times 2 + \frac{1}{8} \times 4 + \frac{1}{16} \times 8 \cdots = 1/2 + 1/2 + 1/2 + \cdots = \infty.$$

   It is a pretty good game for you to play; your expected profit is infinite. This is not really a paradox, but it is counter-intuitive that such a simple game could lead to infinite expected value. One consequence is that repeatedly playing this game and adding the winnings does *not* approach a normal distribution, in contrast to the Central Limit Theorem.

4.
   - If you give $1/3$ probability to all three actions, then you are in good shape: no matter what your opponent does you have $1/3$ chance of winning, losing, or tying, so your expected value is $1/3(1) + 1/3(0) + 1/3(-1) = 0$.
   - To show that this is the *only* winning strategy takes more thought. Let p be the probability you assign to paper, r to rock, and s to scissors. So $p + r + s = 1$ and all three variables are non-negative.

     What values can these variables take on? When your opponent plays rock, you need to have non-negative expected value. This expected value is $r \times 0 + p \times 1 + s \times (-1)$. So it is required that $p - s \geq 0$. This is the same as $p \geq s$.

     Then if you repeat the same calculation for your opponent's other two actions, you get $s - r \geq 0$ and $r - p \geq 0$, so $s \geq r$ and $r \geq p$. But the only way all of $p \geq s \geq r \geq p$ can be true is if all three probabilities are equal!
   - This time we have $p = 0$ and need to pick non-negative r and s with $r + s = 1$. As we just calculated, we want $\min(p - s, s - r, r - p)$ to be as large as possible, which is the same as $\min(-s, s - r, r)$ because of our allergy. If you substitute $s = 1 - r$ this is the same as maximizing $\min(r - 1, 1 - 2r, r)$. Clearly $r - 1 < r$ so we want to maximize $\min(r - 1, 1 - 2r)$. If you draw a graph of these functions, you see that the maximum over all possible choices $0 \leq r \leq 1$ occurs at the intersection where $r - 1 = 1 - 2r$, or $r = 2/3$. So the unique optimal strategy is $p = 0, r = 2/3, s = 1/3$, and since $\min(p - s, s - r, r - p) = -1/3$ you expect to lose at most \$1/3 per round.

5. There are a lot of ways to solve this problem. Here is a solution with a medium amount of work. First, the probability that your symbol comes up 3 times ($T = 3$) is $(1/6)^3 = 1/216$.

Next, the probability that your symbol comes up 0 times ($T = 3$) is $(5/6)^3 = 125/216$, since each die misses your symbol with probability $5/6$. What about getting exactly 1 or 2 symbols? If your symbol comes up exactly once, it can either come up on the first, second, or third die. The probability that

the first die is your symbol and the other two are not

is $1/6 \times 5/6 \times 5/6$. So taking into account the other two possible ways for the symbol to appear exactly once, the overall probability of $T = 1$ is $3(1/6)(5/6)^2 = 75/216$. Likewise, the number of ways for the symbol to appear exactly twice ($T = 2$) is $3(5/6)(1/6)^2 = 15/216$. Then your expected profit is

$$\frac{1}{216} \times 3 + \frac{15}{216} \times 2 + \frac{75}{216} \times 1 + \frac{125}{216} \times (-1) = \frac{3 + 30 + 75 - 125}{216} = \frac{-17}{216}.$$

A cute, but trickier solution: your profit equals $T - Z$ where Z is a random variable equal to 1 when your symbol shows on zero dice (i.e., when $T = 0$) and 0 otherwise. By linearity of expectation your profit is $\mathsf{E}[T - Z] = \mathsf{E}[T] - \mathsf{E}[Z]$. Again by linearity of expectation and the fact that you expect your number to be on top of a die 1/6 of the time, you have $\mathsf{E}[T] = 3/6 = 1/2$. And $\mathsf{E}[Z] = 125/216$. So the answer is $1/2 - 125/216 = -17/216$.

6. We first expand the binomial, and then use linearity of expectation. You then need to replace $\mathsf{E}[X\mathsf{E}[X]]$ by $\mathsf{E}[X]\mathsf{E}[X]$, which is valid because $\mathsf{E}[X]$ is a constant.

$$\mathsf{E}[(X - \mathsf{E}[X])^2] = \mathsf{E}[X^2 - 2X\mathsf{E}[X] + \mathsf{E}[(\mathsf{E}[X])^2] = \mathsf{E}[X^2] - 2\mathsf{E}[X\mathsf{E}[X]] + (\mathsf{E}[X])^2$$
$$= \mathsf{E}[X^2] - 2\mathsf{E}[X]\mathsf{E}[X] + (\mathsf{E}[X])^2 = \mathsf{E}[X^2] - (\mathsf{E}[X])^2.$$

For the second half, because $(X - \mathsf{E}[X])^2$ is never negative, the first definition $\mathsf{E}(X - \mathsf{E}[X])$ of variance is never negative. Thus the second definition of variance is never negative: $\mathsf{E}[X^2] - (\mathsf{E}[X])^2 \geq 0$ giving $\mathsf{E}[X^2] \geq (\mathsf{E}[X])^2$ as we wanted.

7. Watch the video! Both are equally likely at any given point in time (you expect to see $(N - 2)/8$ of either one after flipping N coins). But the informal difference is that (H, T, H) can overlap itself, and clustering of "hits" means also large stretches of "misses." You have to wait longer for (H, T, H) on average.

8. To verify that virtual valuation is preserved by a coin flip, let's say we have \$$x$ in real dollars right now — our current virtual value is $((1-p)/p)^x$. When we flip, we have probability p of going to $x+1$ and probability $1-p$ of going to $x-1$. So the expected virtual value after we flip is

$$p((1-p)/p)^{x+1} + (1-p)((1-p)/p)^{x-1} = \frac{(1-p)^{x+1}}{p^x} + \frac{(1-p)^x}{p^{x-1}} = [(1-p)/p + 1]\frac{(1-p)^x}{p^{x-1}}$$

$$= [1/p]\frac{(1-p)^x}{p^{x-1}} = ((1-p)/p)^x.$$

This is indeed the same as the virtual value before the flip, which was what we wanted to check. To complete the problem, we let s be the probability of success, and note that the initial virtual value equals the final expected virtual value, so

$$((1-p)/p)^{30} = s((1-p)/p)^{100} + (1-s)((1-p)/p)^0.$$

When we solve this equation for s we get the probability mentioned in the problem.