Random Walks

1 Gambler's Ruin

Today we're going to talk about one-dimensional random walks. In particular, we're going to cover a classic phenomenon known as gambler's ruin. The gambler's ruin problem is a particularly good way to end the term since its solution requires several of the techniques that we learned during the term. Those of you who like to gamble are sure to find it interesting.

Suppose we start with n dollars, and make a sequence of bets. For each bet, we win 1 dollar with probability p, and lose 1 dollar with probability $1-p$. We quit if either we go broke, in which case we lose, or when we reach $T = n + m$ dollars, that is, when we win m dollars. For example, in Roulette, $p = \frac{18}{38} = \frac{9}{19} \approx .473$. If $n = 100$ dollars, and $m = 100$ dollars, then $T = 200$ dollars. What are the odds we win 100 dollars before losing 100 dollars? Most folks would think that since $.473 \approx .5$, the odds are not so bad. In fact, as we will see, we win before we lose with probability at most $\frac{1}{37648}$.

This is an amazing result! The classic strategy is to continue betting until you are a little bit ahead, and then quit. We're going to see in just a moment that this strategy is not very good. Even if you quit when you're up 20%, say by making 20 dollars in the above exmaple, we will show the chance of winning before losing is at most $1/8$.

This problem is a classic example of a problem that involves a one-dimensional random walk. In such a random walk, there is some value - say the number of dollars we have - that can go up or down or stay the same at each step with some probabilities. In this example, we have a random walk in which the value can go up or down by 1 at each step. We can diagram it as follows.

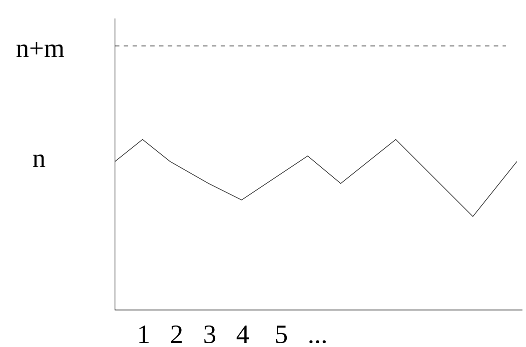

The probability of making an up move at any step is p, no matter what has happened in the past. The probability of making a down move is $1 - p$. This random walk is a special type of random walk where moves are independent of the past, and is called a martingale. If $p = 1/2$, the random walk is unbiased, whereas if $p \neq 1/2$, the random walk is biased. We also have boundaries at 0 and $n + m$. If the walk hits a boundary, then we stop playing, i.e., we quit when broke (lose n) or when we get to $n+m$ (win m). So we care about the probability of winning or the probability of going broke. Note that these do not necessarily sum up to 1 since there is some chance that we never hit a boundary and walk forever. However, we will show later that this chance is essentially 0.

So let's figure out the probability that we gain m before losing n. To set things up formally, let W be the event we hit T before we hit 0, where $T = n + m$. Let D_t be a random variable that denotes the number of dollars we have at time step t. Let $P_n = \Pr(W \mid D_0 = n)$ be the probability we get T before we go broke, given that we start with n dollars. Our question then, is what is P_n?

We're going to use a recursive approach.

Claim 1. $P_n = \begin{cases} 0 & \text{if } n = 0 \\ 1 & \text{if } n = T \\ pP_{n+1} + (1-p)P_{n-1}, & \text{if } 0 < n < T \end{cases}$

The intuition here is clear if $n = 0$ or $n = T$, since in this case we have either already lost or already won. Otherwise, since our moves are independent of the past, with probability p we obtain $n + 1$ dollars, and with probability $1 - p$ we obtain $n - 1$ dollars, which gives the recurrence.

Proof. $P_0 = \Pr(W \mid D_0 = 0) = 0$ since we've already lost, and $P_T = \Pr(W \mid D_0 = T) = 1$ since we've already won. Now, assume $0 < n < T$. Let E be the event that the first bet is a win, and $\bar{E}$ the event that the first bet is a loss. Then,

$$\begin{aligned} P_n &= \Pr(W \mid D_0 = n) \\ &= \Pr(W \wedge E \mid D_0 = n) + \Pr\left(W \wedge \bar{E} \mid D_0 = n\right) \\ &= \Pr(E \mid D_0 = n) \Pr(W \mid E \wedge D_0 = n) + \Pr\left(\bar{E} \mid D_0 = n\right) \Pr\left(W \mid \bar{E} \wedge D_0 = n\right) \\ &= p \Pr(W \mid D_1 = n+1) + (1-p) \Pr(W \mid D_1 = n-1) \\ &= p \Pr(W \mid D_0 = n+1) + (1-p) \Pr(W \mid D_0 = n-1) \\ &= pP_{n+1} + (1-p)P_{n-1}. \end{aligned}$$

The first equation is by definition. The second is by the theorem of total probability. The third is by the definition of conditional probability. The fourth and fifth follow from the fact that E is independent of the current dollar amount, and the time step we are at. □

Now we have a recurrence to solve. Rewriting it, we obtain $pP_{n+1} - P_n + (1-p)P_{n-1} = 0$, where $P_0 = 0$ and $P_T = 1$. This is a linear homogeneous recurrence. We start by solving the characteristic equation $pr^2 - r + (1-p) = 0$. Using the quadratic formula,

$$\begin{aligned} r &= \frac{1 \pm \sqrt{1-4p(1-p)}}{2p} \\ &= \frac{1 \pm \sqrt{1-4p+4p^2}}{2p} \\ &= \frac{1 \pm (1-2p)}{2p} \\ &= \frac{2-2p}{2p},\ \frac{2p}{2p} \\ &= \frac{1-p}{p},\ 1. \end{aligned}$$

Thus, we get two distinct roots iff $p \neq 1/2$. So, in this case we know

$$P_n = A\left(\frac{1-p}{p}\right)^n + B(1)^n = A\left(\frac{1-p}{p}\right)^n + B.$$

Using the boundary conditions, $0 = P_0 = A + B$, so $B = -A$. Also,

$$1 = P_T = A\left(\frac{1-p}{p}\right)^T + B = A\left(\left(\frac{1-p}{p}\right)^T - 1\right),$$

and so

$$A = \frac{1}{\left(\frac{1-p}{p}\right)^T - 1},\ \ B = \frac{-1}{\left(\frac{1-p}{p}\right)^T - 1}.$$

Thus,

$$\begin{aligned} P_n &= \frac{\left(\frac{1-p}{p}\right)^n}{\left(\frac{1-p}{p}\right)^T - 1} - \frac{1}{\left(\frac{1-p}{p}\right)^T - 1} \\ &= \frac{\left(\frac{1-p}{p}\right)^n - 1}{\left(\frac{1-p}{p}\right)^T - 1} \\ &\leq \frac{\left(\frac{1-p}{p}\right)^n}{\left(\frac{1-p}{p}\right)^T} \\ &= \left(\frac{p}{1-p}\right)^{T-n} \\ &= \left(\frac{p}{1-p}\right)^m. \end{aligned}$$

Note that the last expression is even independent of n. It is also exponentially small in m. If $p = 9/19$ in our earlier example, then $p/(1-p) = 9/10$, and for any n, if $m = 100$ dollars, then

$$\Pr(\text{Win}) \leq \left(\frac{9}{10}\right)^{100} \leq \frac{1}{37648}.$$

Thus, even if we start with $n = 10^9$ dollars, we have the same very small upper bound on the probability that we win. Note that if $m = 1000$ dollars, then

$$\Pr(\text{Win}) \leq \left(\frac{9}{10}\right)^{1000} \leq 2 \cdot 10^{-46},$$

so there is really no chance. On the other hand, suppose $n = 10$ dollars and $T = 20$ dollars. In this case the probability of winning is

$$\frac{\left(\frac{10}{9}\right)^{10} - 1}{\left(\frac{10}{9}\right)^{20} - 1} \approx .26,$$

so the odds are better, though not good. So if you're going to gamble, learn to count cards and play blackjack, or bet it all at once!

What's the intuition for this? Normally we would think that the probability of winning 100 dollars before losing 200 dollars is better than winning 10 before losing 10, i.e., that the ratio is what matters. In fact, the ratio *is* what matters if the game is fair, i.e., if $p = 1/2$. Let's look at that case.

When $p = 1/2$, we have the characteristic equation $pr^2 - r + (1-p) = 0$, or $\frac{1}{2}r^2 - r + \frac{1}{2} = 0$, or $r^2 - 2r + 1 = 0$, or $(r-1)^2 = 0$. Thus, there is a double root at $r = 1$.

Thus, $P_n = An(1)^n + B(1)^n = An + B$. Using the boundary conditions, $0 = P_0 = B$ implies $B = 0$, and $1 = P_T = AT + B = AT$ implies $A = \frac{1}{T}$. Thus, for $p = \frac{1}{2}$, we obtain

$$P_n = \frac{n}{T} = \frac{n}{n+m}.$$

So the probability of getting to T before hitting 0 is $\frac{n}{T}$. This is closer to what our intuition expects. So the trouble comes when the game is not fair. Indeed, in this case, if $n = 200$ and $m = 100$ (so $T = 300$), we have $\Pr(\text{Win}) = \frac{200}{300} = \frac{2}{3}$. On the other hand, if $n = 10$ and $m = 10$ (so $T = 20$), then $\Pr(\text{Win}) = \frac{1}{2}$. Thus, actually, now we are more likely to win in the first case!

So the trouble is that our intuition tells us that if the game is almost fair, then we expect the results to be almost the same as if the game were fair. It turns out this is not the case!

For the unbiased case, the intuition here is that it takes a long time to hit the boundary, but we will do so eventually (we'll talk about time later), and the walk is symmetric about the starting amount.

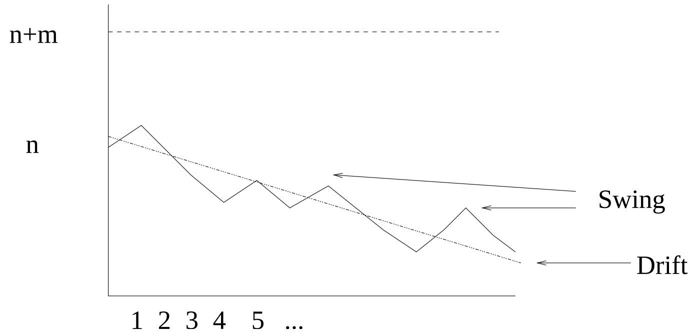

On the other hand, for the biased case, by the time you are likely to swing high, your baseline has already dropped too low for the swing to hit T. The downward drift seems small, but it really does dominate the effect of the swings. So the difference between the unbiased and biased games is the drift. The swings are similar but the drift takes over in biased games.

2 The time to completion

So we've figured out the probability of winning some number of dollars before going broke. As we saw, this probability is very low for biased games. It's logical to conclude that this means that we're very likely to end the game without any money.

Before we can conclude this, though, we need to rule out one other possibility, namely, the possibility that we play forever. Indeed, it is possible that you can play forever without winning the required amount and without going broke. In fact, it turns out that the probability that you never end is 0 even for unbiased games!

Theorem 2. *$\forall p$, the probability that the game never ends is 0.*

This seems paradoxical since there are lots of walks that go forever, yet the probability of going forever is still zero. To really understand this, we need to get into measure theory, which will say that the set of non-terminating walks is an infinite set of measure 0. In any case, we can prove the theorem now.

Proof. We can assume $0 < p < 1$, since if $p = 0$ we go broke and if $p = 1$ we win, so the game ends. Now define $Q_n = \Pr(\text{Play forever} \mid D_0 = n)$. We first claim that $Q_n = \begin{cases} 0 & \text{if } n = 0 \\ 0 & \text{if } n = T \\ pQ_{n+1} + (1-p)Q_{n-1}, & \text{if } 0 < n < T \end{cases}$

This is the same as our constraint on P_n, except now we have $Q_n = 0$ if $n = T$ (rather than the $P_T = 1$ we had before). Hence, Q_n has the same recurrence as P_n, except the boundary condition is different. Thus,

$$Q_n = A\left(\frac{1-p}{p}\right)^n + B \quad \text{if } p \neq \frac{1}{2}$$

We can solve for A and B with the new boundary conditions, obtaining $0 = A + B$, so $B = -A$, and $0 = A\left(\frac{1-p}{p}\right)^T + B$, so $A\left(\frac{1-p}{p}\right)^T - A = 0$. But for $p \neq \frac{1}{2}$, this implies $A = 0$. Thus, $A = B = 0$, so $Q_n = 0$, so the probability we play forever is 0.

Now we check the case when $p = \frac{1}{2}$. In this case $Q_n = An + B$. Using the boundary conditions, $0 = Q_0 = B$, so $B = 0$. Also, $0 = Q_T = AT + B$, so $A = 0$. Thus, $Q_n = 0$, as desired. □

The preceding argument tells us that the probability of playing forever is 0, but it doesn't give a very good idea of how long you should expect to play. It turns out that you can also compute the expected time to win or go broke by using recurrences.

Definition 1. *Let S be the number of steps until we hit a boundary, that is, until we win or lose. Let $E_n = \text{Ex}\,(S \mid D_0 = n)$ be the expected time to win or lose given that we start with n dollars.*

Claim 3. $E_n = \begin{cases} 0 & \text{if } n = 0 \\ 0 & \text{if } n = T \\ 1 + pE_{n+1} + (1-p)E_{n-1}, & \text{if } 0 < n < T \end{cases}$

Proof. The case when $n = 0$ or T is immediate, since we have already lost or won. The case when $0 < n < T$ is similar to before. We have,

$$\begin{aligned} E_n &= \text{Ex}\,(S \mid D_0 = n) \\ &= \text{Ex}\,(S \mid D_0 = n \wedge \text{ win 1st bet})\,\Pr\,(\text{win 1st bet}) \\ &+ \text{Ex}\,(S \mid D_0 = n \wedge \text{ lose 1st bet})\,\Pr\,(\text{lose 1st bet}) \\ &= \text{Ex}\,(S \mid D_1 = n+1)\,p + \text{Ex}\,(S \mid D_n = n-1)\,(1-p) \\ &= (1 + \text{Ex}\,(S \mid D_0 = n+1))p + (1 + \text{Ex}\,(S \mid D_0 = n-1))(1-p) \\ &= p + pE_{n+1} + 1 - p + (1-p)E_{n-1} \\ &= 1 + pE_{n+1} + (1-p)E_{n-1}. \end{aligned}$$

□

So now we just need to solve the recurrence for E_n. Rewriting, we have

$$pE_{n+1} - E_n + (1-p)E_{n-1} = -1,$$

with $E_0 = 0$ and $E_T = 0$. The difference here is just that the recurrence is inhomogeneous. First, we take a homogeneous solution,

$$E_n = A\left(\frac{1-p}{p}\right)^n + B \quad \text{if } p \neq \frac{1}{2},$$

which is the same as before. Next, we guess a particular solution $E_n = an + b$. Plugging in the particular solution,

$$\begin{aligned} p(a(n+1)+b) - (an+b) + (1-p)(a(n-1)+b) &= -1 \\ pan + pa + pb - an - b + an - a + b - pan + pa - pb &= -1 \\ 2pa - a &= -1 \\ a(2p-1) &= -1 \\ a &= \frac{1}{1-2p}, \end{aligned}$$

and b is unconstrained. So we can set $b = 0$. Thus,

$$E_n = A\left(\frac{1-p}{p}\right)^n + B + \frac{n}{1-2p}.$$

We now solve for A and B using that $E_0 = E_T = 0$. We have $0 = E_0 = A + B$, so $B = -A$. We have $0 = E_T = A\left(\frac{1-p}{p}\right)^T + B + \frac{T}{1-2p}$, and thus $A = \frac{-\frac{T}{(1-2p)}}{\left(\frac{1-p}{p}\right)^T - 1}$. Plugging in A and B, and simplifying, we get

$$E_n = \frac{n}{1-2p} - \frac{T}{(1-2p)} \cdot \frac{\left(\frac{1-p}{p}\right)^n - 1}{\left(\frac{1-p}{p}\right)^T - 1} \quad \text{if } p \neq \frac{1}{2}$$

This expression looks messy, but we can conclude a few nice things. First, for $p < 1/2$, $E_n \leq \frac{n}{1-2p}$, and as $T \to \infty$, $E_n \to \frac{n}{1-2p}$. This is not surprising since we expect to lose $1-2p$ dollars each bet. Indeed, the expected loss per bet is $(1-p) - p = 1 - 2p$. If we actually were to lose $1-2p$ dollars each bet, then we would go broke in exactly $\frac{n}{1-2p}$ steps.

We must also consider the case when $p = 1/2$. In that case we have the homogeneous solution $E_n = An + B$. If we were to guess $E_n = an + b$ as a particular solution, we'd have

$$\begin{aligned} -1 &= pE_{n+1} - E_n + (1-p)E_{n-1} \\ &= p(an + a + b) - an - b + (1-p)(an - a + b) \\ &= pan + pa + pb - an - b + an - a + b - pan + pa - pb \\ &= 2pa - a \\ &= 2 \cdot \frac{1}{2}a - a \\ &= 0, \end{aligned}$$

which is clearly a contradiction. Due to the repeated root, we need to guess a higher-order polynomial. We guess $E_n = an^2 + bn + c$. Plugging this into the recurrence, some algebra shows that it works with $a = -1$ and $b = c = 0$. Thus,

$$E_n = An + B - n^2.$$

Now for the boundary conditions. We have $E_0 = 0 = B$, so $B = 0$. Also, $E_T = 0 = AT - T^2$, so $A = T$. Thus, $E_n = Tn - n^2$. But since $T = n + m$, we have $E_n = (n+m)n - n^2$, or

$$E_n = nm \quad \text{if } p = \frac{1}{2}.$$

This is a very clean result, which states that in an unbiased game, you expect to play for the product of the amount you're willing to lose times the amount you want to win. As a corollary, if you never quit when you're ahead, then you can expect to play forever in a fair game. That is, as $m \to \infty$, $E_n \to \infty$. So if you play until you go broke, you can expect to play forever. This is very good news if you like to gamble.

We can prove this by observing that the expected time to go broke is at least the expected time to go broke or hit m for all m. Let E_n^* be the expected time to go broke starting at n. Then for all m, $E_n^* \geq nm$, which means that $E_n^* = \infty$ for $n > 0$ and $p = \frac{1}{2}$.

This holds even for $n = 1$ dollar. Thus, starting with just 1 dollar, you still expect to play forever! Well, not exactly, more precisely the expected time to go broke is ∞. Note that you could easily go broke, and 50% of the time this happens on the first bet. In fact, even though you expect to play an infinite number of bets, the probability you will eventually go broke is 1, as we will now show. This seems impossible, but things like this can happen when the expectation is infinite.

Theorem 4. *If you start with n dollars and $p = \frac{1}{2}$ and you play until you go broke, then for all n,* $\Pr(\textit{go broke}) = 1$.

This says that even if you start with 1 million dollars and are playing an unbiased (i.e., fair) game, with probability 1 you will eventually lose it all.

Proof. For all T,

$$\begin{aligned}
\Pr(\text{go broke}) &\geq \Pr(\text{go broke before reach } T) \\
&= 1 - \Pr(\text{reach } T \text{ before go broke}) \quad (\text{since } \Pr(\text{ play forever }) = 0) \\
&= 1 - \frac{n}{T} \\
&\to 1 \text{ as } T \to \infty
\end{aligned}$$

In other words, the probability you go broke cannot be less than 1, so it must be 1. □

So even if you play a fair game and choose not to quit, you will eventually go broke.

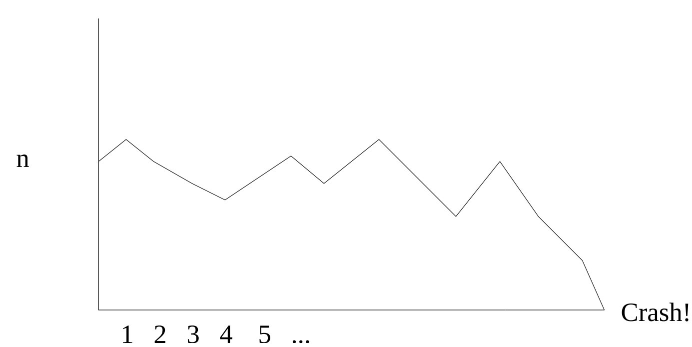

The result is not true if $p > 1/2$. In this case there is a non-zero chance that you can play forever and never go broke. We'll cover this next year.

3 Appendix

In lecture the following question concerning a generalization of the gambler's ruin problem was asked. Suppose there is a finite sequence $z_1, z_2, \ldots, z_s$ of distinct integers, such that each $z_i \geq -1$. Let $p_1, p_2, \ldots, p_s$ be a sequence of real numbers with $0 < p_i < 1$ for all i, and $\sum_{i=1}^{s} p_i = 1$. Consider a variation of gambler's ruin in which at each step, the amount of money the gambler has changes by z_i (which may be negative) with probability p_i. Define the *downward drift* d to be

$$d = -\sum_{i=1}^{s} p_i z_i.$$

Suppose you start with $n > 1$ dollars. The question is what the expected time E_n is until you end up with 0 dollars, assuming you never quit otherwise. If no z_i is equal to -1, then since the z_i are integers, they are all non-negative. Clearly in this case $E_n = \infty$ since you will never end up with less than n dollars. Assume, then, w.l.o.g. that $z_1 = -1$.

By our earlier arguments, we know

$$E_1 = 1 + \sum_{i>1} p_i E_{1+z_i}.$$

Consider E_n for any integer $n > 1$. Let R_i be the number of steps in between the first time you have i dollars and the first time you have $i-1$ dollars. By linearity of expectations, $E_n = \sum_{i=1}^{n} \mathrm{Ex}\left(R_i\right)$. By symmetry, $\mathrm{Ex}\left(R_i\right) = E_1$ for all i. Thus, $E_n = nE_1$. Thus,

$$E_1 = 1 + \sum_{i>1} p_i (1 + z_i) E_1,$$

and thus either E_1 is infinite or

$$\begin{aligned}
E_1 &= \frac{1}{1 - \sum_{i>1} p_i (1 + z_i)} \\
&= \frac{1}{1 - \sum_{i>1} p_i - \sum_{i>1} p_i z_i} \\
&\geq \frac{1}{1 - (1 - p_1) - (-d - p_1(-1))} \\
&= \frac{1}{1 - 1 + p_1 - p_1 + d} \\
&= \frac{1}{d}.
\end{aligned}$$

Thus, either E_n is infinite, or is $\frac{n}{d}$. Let P_i be the probability we don't go broke until at least i steps. Now, E_n is just

$$\sum_{i=1}^{\infty} P_i.$$

One can show using Chernoff bounds that for a sufficiently large integer i_0, for all $i > i_0$

$$P_i \leq c^{-i},$$

for some constant $c > 1$. Thus, we can upper bound the above sum as,

$$\sum_{i=0}^{i_0} P_i + \sum_{i_0+1}^{\infty} c^{-i}.$$

The first sum is over a finite number of real numbers, and is thus finite, and the second sum is a geometric series, and is thus finite. Thus, E_n is finite.

The result, however, does not hold if you allow some of the z_i to be less than -1. For instance, if $z_1 = -1, z_2 = -2, p_1 = p$, and $p_2 = 1 - p$, then $E_1 = 1$, $E_2 = 1 + p$, and $E_3 = 2 + p^2$. In this case the above would say that $E_n = n/(2 - p)$, but this is impossible unless $p = 1$.