

20 Random Walks

Random Walks are used to model situations in which an object moves in a sequence of steps in randomly chosen directions. Many phenomena can be modeled as a random walk and we will see several examples in this chapter. Among other things, we’ll see why it is rare that you leave the casino with more money than you entered with and we’ll see how the Google search engine uses random walks through the graph of the world-wide web links to determine the relative importance of websites.


20.1 Unbiased Random Walks

20.1.1 A Bug’s Life

There is a small flea named Stencil. To his right, there is an endless flat plateau. One inch to his left is the Cliff of Doom, which drops to a raging sea filled with flea-eating monsters.



Each second, Stencil hops 1 inch to the right or 1 inch to the left with equal probability, independent of the direction of all previous hops. If he ever lands on the very edge of the cliff, then he teeters over and falls into the sea. So, for example, if Stencil’s first hop is to the left, he’s fishbait. On the other hand, if his first few hops are to the right, then he may bounce around happily on the plateau for quite



some time.

Our job is to analyze the life of Stencil. Does he have any chance of avoiding a fatal plunge? If not, how long will he hop around before he takes the plunge?

Stencil’s movement is an example of a *random walk*. A typical *one-dimensional* random walk involves some value that randomly wavers up and down over time. The walk is said to be *unbiased* if the value is equally likely to move up or down. If the walk ends when a certain value is reached, then that value is called a *boundary condition* or *absorbing barrier*. For example, the Cliff of Doom is a boundary condition in the example above.

Many natural phenomena are nicely modeled by random walks. However, for some reason, they are traditionally discussed in the context of some social vice. For example, the value is often regarded as the position of a drunkard who randomly staggers left, staggers right, or just wobbles in place during each time step. Or the value is the wealth of a gambler who is continually winning and losing bets. So discussing random walks in terms of fleas is actually sort of elevating the discourse.

20.1.2 A Simpler Problem

Let’s begin with a simpler problem. Suppose that Stencil is on a small island; now, not only is the Cliff of Doom 1 inch to his left, but also there is another boundary condition, the Pit of Disaster, 2 inches to his right! For example, see Figure 20.1

In the figure, we’ve worked out a tree diagram for Stencil’s possible fates. In




Figure 20.1 An unbiased, one-dimensional random walk with absorbing barriers at positions 0 and 3. The walk begins at position 1. The tree diagram shows the probabilities of hitting each barrier.

particular, he falls off the Cliff of Doom on the left side with probability:

$$\begin{aligned} \frac{1}{2} + \frac{1}{8} + \frac{1}{32} + \dots &= \frac{1}{2} \left(1 + \frac{1}{4} + \frac{1}{16} + \dots \right) \left(\frac{1}{2} \right) \\ &= \frac{1}{2} \cdot \frac{1}{1 - 1/4} \\ &= \frac{2}{3}. \end{aligned}$$


Similarly, he falls into the Pit of Disaster on the right side with probability:

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

There is a remaining possibility: Stencil *could* hop back and forth in the middle of the island forever. However, we’ve already identified two disjoint events with probabilities 2/3 and 1/3, so this happy alternative must have probability 0.

20.1.3 A Big Island

Putting Stencil on such a tiny island was sort of cruel. Sure, he’s probably carrying bubonic plague, but there’s no reason to pick on the little fella. So suppose that we instead place him n inches from the left side of an island w inches across: In



other words, Stencil starts at position n and his random walk ends if he ever reaches positions 0 or w .

Now he has three possible fates: he could fall off the Cliff of Doom, fall into the Pit of Disaster, or hop around on the island forever. We could compute the probabilities of these three events with a horrific summation, but fortunately there’s a far easier method: we can use a linear recurrence.

Let R_n be the probability that Stencil falls to the right into the Pit of Disaster, given that he starts at position n . In a couple special cases, the value of R_n is easy to determine. If he starts at position w , he falls into the Pit of Disaster immediately, so $R_w = 1$. On the other hand, if he starts at position 0 , then he falls from the Cliff of Doom immediately, so $R_0 = 0$.

Now suppose that our frolicking friend starts somewhere in the middle of the island; that is, $0 < n < w$. Then we can break the analysis of his fate into two cases based on the direction of his first hop:

- If his first hop is to the left, then he lands at position $n - 1$ and eventually falls into the Pit of Disaster with probability R_{n-1} .
- On the other hand, if his first hop is to the right, then he lands at position $n + 1$ and eventually falls into the Pit of Disaster with probability R_{n+1} .

Therefore, by the Total Probability Theorem, we have:

$$R_n = \frac{1}{2}R_{n-1} + \frac{1}{2}R_{n+1}.$$

Solving the Recurrence

Let’s assemble all our observations about R_n , the probability that Stencil falls into the Pit of Disaster if he starts at position n :

$$\begin{aligned} R_0 &= 0 \\ R_w &= 1 \\ R_n &= \frac{1}{2}R_{n-1} + \frac{1}{2}R_{n+1} \quad (0 < n < w). \end{aligned}$$

This is just a linear recurrence—and we know how to solve those! Uh, right? Remember Chapter 10 or Chapter 12?

There is one unusual complication: in a normal recurrence, R_n is written a function of preceding terms. In this recurrence equation, however, R_n is a function of both a preceding term (R_{n-1}) and a *following* term (R_{n+1}). This is no big deal, however, since we can just rearrange the terms in the recurrence equation:

$$R_{n+1} = 2R_n - R_{n-1}.$$

Now we’re back on familiar territory.

Let’s solve the recurrence. The characteristic equation is:

$$x^2 - 2x + 1 = 0.$$

This equation has a double root at $x = 1$. There is no inhomogeneous part, so the general solution has the form:

$$R_n = a \cdot 1^n + b \cdot n1^n = a + bn.$$

Substituting in the boundary conditions $R_0 = 0$ and $R_w = 1$ gives two linear equations:

$$\begin{aligned} 0 &= a, \\ 1 &= a + bw. \end{aligned}$$

The solution to this system is $a = 0$, $b = 1/w$. Therefore, the solution to the recurrence is:

$$R_n = n/w.$$

20.1.4 Death Is Certain

Our analysis shows that if we place Stencil n inches from the left side of an island w inches across, then he falls off the right side with probability n/w . For example, if Stencil is $n = 4$ inches from the left side of an island $w = 12$ inches across, then he falls off the right side with probability $n/w = 4/12 = 1/3$.

We can compute the probability that he falls off the *left* side by exploiting the symmetry of the problem: the probability that he falls off the *left* side starting at position n is the same as the probability that he falls off the *right* side starting at position $w - n$, which is $(w - n)/w$.

This is bad news. The probability that Stencil eventually falls off one side or the other is:

$$\frac{n}{w} + \frac{w - n}{w} = 1.$$

There’s no hope! The probability that Stencil hops around on the island forever is zero.

And there’s even worse news. Let’s go back to the original problem where Stencil is 1 inch from the left edge of an *infinite* plateau. In this case, the probability that he eventually falls into the sea is:

$$\lim_{w \rightarrow \infty} \frac{w - 1}{w} = 1.$$

So even if there were no Pit of Disaster, Stencil still falls off the Cliff of Doom with probability 1. And since


$$\lim_{w \rightarrow \infty} \frac{w - n}{w} = 1$$

for any finite n , this is true no matter where Stencil starts. Our little friend is doomed!

Hey, you know how in the movies they often make it look like the hero dies, but then he comes back in the end and everything turns out okay? Well, we’re not sayin’ anything, just pointing that out.

20.1.5 Life Expectancy

On the bright side, Stencil may get to hop around for a while before he goes over an edge. Let’s use the same setup as before, where he starts out n inches from the left side of an island w inches across: What is the expected number of hops he takes



before falling off an edge?

Let X_n be Stencil’s expected lifespan, measured in hops. If he starts at either edge of the island, then he dies immediately:

$$\begin{aligned} X_0 &= 0, \\ X_w &= 0. \end{aligned}$$

If he starts somewhere in the middle of the island ($0 < n < w$), then we can again break down the analysis into two cases based on his first hop:

- If his first hop is to the left, then he lands at position $n - 1$ and can expect to live for another X_{n-1} steps.
- If his first hop is to the right, then he lands at position $n + 1$ and can expect to live for another X_{n+1} steps.

Thus, by the Law of Total Expectation and Linearity of Expectation, Stencil’s expected lifespan is:

$$X_n = 1 + \frac{1}{2}X_{n-1} + \frac{1}{2}X_{n+1}.$$

The leading 1 accounts for his first hop.

Solving the Recurrence

Once again, Stencil’s fate hinges on a recurrence equation:

$$\begin{aligned} X_0 &= 0 \\ X_w &= 0 \\ X_n &= 1 + \frac{1}{2}X_{n-1} + \frac{1}{2}X_{n+1} \quad (0 < n < w). \end{aligned}$$

We can rewrite the last line as:

$$X_{n+1} = 2X_n - X_{n-1} - 2. \tag{20.1}$$

As before, the characteristic equation is:

$$x^2 - 2x + 1 = 0.$$

There is a double-root at 1, so the homogeneous solution has the form:

$$X_n = a + bn.$$

But this time, there’s an inhomogeneous term, so we also need to find a particular solution. Since this term is a constant, we should try a particular solution of the form $X_n = c$ and then try $X_n = c + dn$ and then $X_n = c + dn + en^2$ and so forth. As it turns out, the first two possibilities don’t work, but the third does. Substituting $X_n = c + dn + en^2$ into Equation 20.1 gives

$$c + d(n + 1) + e(n + 1)^2 = 2(c + dn + en^2) - (c + d(n - 1) + e(n - 1)^2) - 2,$$

which simplifies to $e = -1$. Since all the c and d terms cancel, $X_n = c + dn - n^2$ is a particular solution for all c and d . For simplicity, let’s take $c = d = 0$. Thus, our particular solution is $X_n = -n^2$.

Adding the homogeneous and particular solutions gives the general form of the solution:

$$X_n = a + bn - n^2.$$

Substituting in the boundary conditions $X_0 = 0$ and $X_w = 0$ gives two linear equations:

$$\begin{aligned} 0 &= a, \\ 0 &= a + bw - w^2. \end{aligned}$$

The solution to this system is $a = 0$ and $b = w$. Therefore, the solution to the recurrence equation is:

$$X_n = wn - n^2 = n(w - n).$$

Interpreting the Solution

Stencil’s expected lifespan is $X_n = n(w - n)$, which is the *product* of the distances to the two edges. Thus, for example, if he’s 4 inches from the left edge and 8 inches from the right cliff, then his expected lifespan is $4 \cdot 8 = 32$.

Let’s return to the original problem where Stencil has the Cliff of Doom 1 inch to his left and an infinite plateau to his right. (Also, cue the “hero returns” theme music.) In this case, his expected lifespan is:

$$\lim_{w \rightarrow \infty} 1(w - 1) = \infty$$

Yes, Stencil is certain to eventually fall off the Cliff of Doom—but his expected lifespan is infinite! This sounds almost like a contradiction, but both answers are correct!

Here’s an informal explanation. It turns out that the probability p_k that Stencil falls from the Cliff of Doom on the k th step is $\Theta(1/k^{3/2})$. You can verify by the integration bound that $\sum_{i=1}^{\infty} 1/k^{3/2}$ converges.

On the other hand, the expected time until Stencil falls over the edge is

$$\begin{aligned} \sum_{i=1}^{\infty} k p_k &\geq c \sum_{k=1}^{\infty} \left(\frac{k}{k^{3/2}} \right) \\ &= c \sum_{k=1}^{\infty} \frac{1}{\sqrt{k}} \\ &= \infty, \end{aligned}$$

where c is a constant that comes from the Θ notation. So our answers are compatible.

20.1.6 Application to Fair Gambling Games

We took the high road for a while, but let’s now discuss random walks in a more conventional setting—gambling.

A gambler goes to Las Vegas with $\$n$ in her pocket. Her plan is to make only $\$1$ bets and somehow she has found a casino that will offer her truly even odds¹; namely, she will win or lose $\$1$ on each bet with probability $1/2$. She’ll play until she is broke or she has won $\$m$. In the latter case, she will go home with

$$w = n + m$$

dollars. What’s the probability that she goes home a winner?

This is identical to the flea problem that we just analyzed. Going broke is analogous to falling off the Cliff of Doom. Going home a winner is analogous to falling into the Pit of Disaster, just a lot more fun.

Our analysis of Stencil’s life tells us everything we want to know about the gambler’s prospects:

- The gambler goes home broke with probability

$$\frac{n}{w} = \frac{m}{n + m},$$

- the gambler goes home a winner with probability

$$\frac{w - n}{w} = \frac{n}{n + m},$$

- the gambler goes home with probability

$$\frac{n}{n + m} + \frac{m}{n + m} = 1,$$

- and the number of bets before the gambler goes home is expected to be

$$n(w - n) = nm.$$

If the gambler gets greedy and keeps playing until she goes broke, then

- the gambler eventually goes broke with probability 1, and
- the number of bets before the gambler goes broke is expected to be infinite.

The bottom line here is clear: when gambling, quit while you are ahead—if you play until you go broke, you will certainly go broke.

And that’s the good news! Matters get much worse for the more typical scenario where the odds are against you. Let’s see why.

¹Don’t worry, we’ll get to the more realistic scenario when she is more likely to lose than win in a moment, but let’s just fantasize about the fair scenario for a bit.

20.2 Gambler’s Ruin

So far, we have considered *unbiased* random walks, where the probability of moving up or down (or left or right) is $1/2$. Unfortunately, things are never quite this simple (or fair) in real casinos.

For example, suppose the gambler goes to Las Vegas and makes \$1 bets on red or black in roulette. In this case, she will win \$1 with probability

$$\frac{18}{38} \approx 0.473$$

and she will lose \$1 with probability

$$\frac{20}{38} \approx 0.527.$$

That’s because the casinos add those bothersome green 0 and 00 to give the house a slight advantage.

At first glance (or after a few drinks), $18/38$ seems awfully close to $1/2$ and so our intuition tells us that the game is “almost fair.” So we might expect the analysis we just did for the fair game to be “almost right” for the real game. For example, if the gambler starts with \$100 and quits when she gets ahead by \$100 in the fair game, then she goes home a winner with probability

$$\frac{100}{200} = .5.$$

And, if she wants to improve her chances of going home a winner, she could bring more money. If she brings \$1000 and quits when she gets ahead by \$100 in the fair game, then she goes home a winner with probability

$$\frac{1000}{1100} \approx .91.$$

So, given that the real game is “almost fair,” we might expect the probabilities of going home a winner in these two scenarios to be “almost 50% and 91%,” respectively.

Unfortunately for the gambler, all this “almost” reasoning will almost surely lead to disaster. Here are the grim facts for the real game where the gambler wins \$1 with probability $18/38$.

n = starting wealth	probability she reaches $n + \$100$ before \$0
\$100	1 in 37649.619496...
\$1000	1 in 37648.619496...
\$1,000,000,000	1 in 37648.619496...

Except on the very low end, the amount of money she brings makes almost no difference!² She is almost certain to go broke before winning \$100. Let’s see why.

20.2.1 Finding a Recurrence

We can approach the gambling problem the same way we studied the life of Stencil. Suppose that the gambler starts with n dollars. She wins each bet with probability p and plays until she either goes bankrupt or has $w = n + m$ dollars in her pocket. (To be clear, w is the total amount of money she wants to end up with, not the amount by which she wants to increase her wealth, which is m .) Our objective is to compute R_n , the probability that she goes home a winner.

As usual, we begin by identifying some boundary conditions. If she starts with no money, then she’s bankrupt immediately so $R_0 = 0$. On the other hand, if she starts with w dollars, then she’s an instant winner, so $R_w = 1$.

Now we divide the analysis of the general situation into two cases based on the outcome of her first bet:

- She wins her first bet with probability p . She then has $n + 1$ dollars and probability R_{n+1} of reaching her goal of w dollars.
- She loses her first bet with probability $1 - p$. This leaves her with $n - 1$ dollars and probability R_{n-1} of reaching her goal.

Plugging these facts into the Total Probability Theorem gives the equation:

$$R_n = pR_{n+1} + (1 - p)R_{n-1}. \quad (20.2)$$

20.2.2 Solving the Recurrence

Rearranging the terms in Equation 20.2 gives us a recurrence for R_n , the probability that the gambler reaches her goal of w dollars if she starts with n :

$$\begin{aligned} R_0 &= 0 \\ R_w &= 1 \\ pR_{n+1} - R_n + (1 - p)R_{n-1} &= 0 \quad (0 < n < w). \end{aligned}$$

The characteristic equation is:

$$px^2 - x + (1 - p) = 0.$$

²The fact that only one digit changes from the first case to the second is a peripheral bit of bizarreness that we’ll leave in your hands.

The quadratic formula gives the roots:

$$\begin{aligned} x &= \frac{1 \pm \sqrt{1 - 4p(1 - p)}}{2p} \\ &= \frac{1 \pm \sqrt{(1 - 2p)^2}}{2p} \\ &= \frac{1 \pm (1 - 2p)}{2p} \\ &= \frac{1 - p}{p} \text{ or } 1. \end{aligned}$$

There’s an important point lurking here. If the gambler is equally likely to win or lose each bet, then $p = 1/2$, and the characteristic equation has a double root at $x = 1$. This is the situation we considered in the flea problem. The double root led to a general solution of the form:

$$R_n = a + bn$$

Now suppose that the gambler is *not* equally likely to win or lose each bet; that is, $p \neq 1/2$. Then the two roots of the characteristic equation are different, which means that the solution has a completely different form:

$$R_n = a \cdot \left(\frac{1 - p}{p}\right)^n + b \cdot 1^n$$

In mathematical terms, this is where the fair game and the “almost fair” game take off in completely different directions: in one case we get a linear solution and in the other we get an exponential solution! This is going to be bad news for anyone playing the “almost fair” game.

Anyway, substituting the boundary conditions into the general form of the solution gives a system of linear equations:

$$\begin{aligned} 0 &= a + b \\ 1 &= a \cdot \left(\frac{1 - p}{p}\right)^w + b. \end{aligned}$$

Solving this system, gives:

$$a = \frac{1}{\left(\frac{1 - p}{p}\right)^w - 1}, \quad b = -\frac{1}{\left(\frac{1 - p}{p}\right)^w - 1}.$$

Substituting these values back into the general solution gives:

$$R_n = \left(\frac{1}{\left(\frac{1-p}{p}\right)^w - 1} \right) \left(\left(\frac{1-p}{p}\right)^n - \frac{1}{\left(\frac{1-p}{p}\right)^w - 1} \right) \\ = \frac{\left(\frac{1-p}{p}\right)^n - 1}{\left(\frac{1-p}{p}\right)^w - 1}.$$

(Suddenly, Stencil’s life doesn’t seem so bad, huh?)

20.2.3 Bad News!

We have an answer! But it’s not good news. If the gambler starts with n dollars and wins each bet with probability p , then the probability she reaches w dollars before going broke is:

$$\frac{\left(\frac{1-p}{p}\right)^n - 1}{\left(\frac{1-p}{p}\right)^w - 1}.$$

Let’s try to make sense of this expression. If the game is biased against her, as with roulette, then $1 - p$ (the probability she loses) is greater than p (the probability she wins). If n , her starting wealth, is also reasonably large, then both exponentiated fractions are big numbers and the -1’s don’t make much difference. Thus, her probability of reaching w dollars is very close to:

$$\left(\frac{1-p}{p}\right)^{n-w} \approx \left(\frac{1-p}{p}\right)^m.$$

In particular, if she is hoping to come out $m = \$100$ ahead in roulette, then $p \approx \frac{18}{38}$ and her probability of success is:

$$\left(\frac{10}{9}\right)^{-100} = 1 \text{ in } 37648.619496.$$

This explains the strange number we arrived at earlier! In fact, this number does not change no matter how large n gets, so even if the gambler starts with a trillion dollars, she is still not likely to ever get ahead by even \$100.

20.2.4 But Why?

Why does the gambler’s starting wealth have so little impact on her probability of coming out ahead? Intuitively, there are two forces at work. First, the gambler’s

wealth has random upward and downward *swings* due to runs of good and bad luck. Second, her wealth has a steady, downward *drift* because she has a small expected loss on every bet. The situation is illustrated in Figure 20.2.

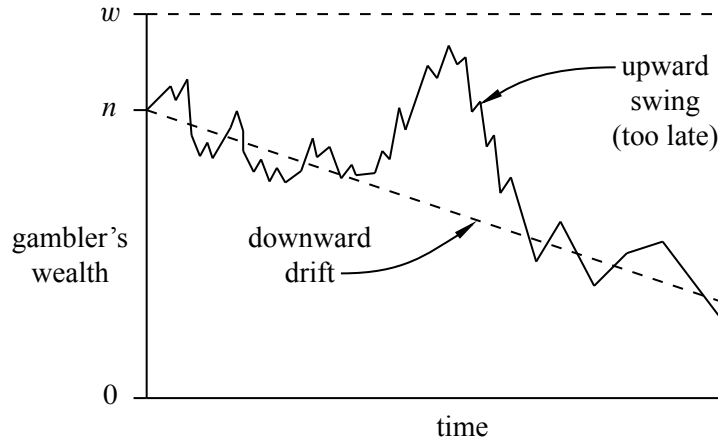


Figure 20.2 In a biased random walk, the downward drift usually dominates swings of good luck.

For example, in roulette, the gambler wins a dollar with probability $18/38$ and loses a dollar with probability $20/38$. Therefore, her expected return on each bet is

$$1 \cdot \frac{18}{38} + (-1) \cdot \frac{20}{38} = \frac{-2}{38} = -\frac{1}{19}.$$

Thus, her expected wealth drifts downward by a little over 5 cents per bet.

One might think that if the gambler starts with a billion dollars, then she will play for a long time, so at some point she should have a lucky, upward swing that puts her \$100 ahead. The problem is that her capital is steadily drifting downward. And after her capital drifts down a few hundred dollars, she needs a huge upward swing to save herself. And such a huge swing is extremely improbable. So if she does not have a lucky, upward swing early on, she’s doomed forever. As a rule of thumb, *drift* dominates *swings* over the long term.

20.2.5 Expected Playing Time

Even though casino gamblers are destined to lose, some of them enjoy the process. So let’s figure out how long their enjoyment is expected to last.

Let X_n be the expected number of bets before going home (broke or a winner).

Reasoning as in Section 20.1.5, we can set up a recurrence for X_n :

$$\begin{aligned} X_0 &= 0, \\ X_w &= 0, \\ X_n &= 1 + (1 - p)X_{n-1} + pX_{n+1}. \end{aligned} \tag{20.3}$$

This is the same as the recurrence for R_n in Equation 20.2 except for the inhomogeneous part.

To find the particular solution, we try $X_n = c$ (which doesn't work) and then $X_n = c + dn$ (which does work as long as $p \neq 1/2$). Plugging $X_n = c + dn$ into Equation 20.3 yields:

$$\begin{aligned} c + dn &= 1 + (1 - p)(c + d(n - 1)) + p(c + d(n + 1)) \\ &= 1 + c + dn - (1 - p)d + pd \end{aligned}$$

and thus that

$$d = \frac{1}{1 - 2p}.$$

Since c is arbitrary, we will set $c = 0$ and our particular solution is

$$X_n = \frac{n}{1 - 2p}.$$

The characteristic equation for Equation 20.3 is

$$px^2 - x + (1 - p) = 0.$$

We have already determined that the roots for this equation are

$$\frac{1 - p}{p} \quad \text{and} \quad 1.$$

Hence, the general solution to the recurrence is

$$X_n = a \left(\frac{1 - p}{p} \right)^n + b + \frac{n}{1 - 2p}.$$

Plugging in the boundary conditions, we find that

$$\begin{aligned} 0 &= a + b, \\ 0 &= a \left(\frac{1 - p}{p} \right)^w + b + \frac{w}{1 - 2p}. \end{aligned}$$

Hence

$$a = \frac{-\left(\frac{w}{1-2p}\right)\left(\frac{1-p}{p}\right)^w}{\left(\frac{1-p}{p}\right)^w - 1} \quad \text{and} \quad b = \frac{\left(\frac{w}{1-2p}\right)\left(\frac{1-p}{p}\right)^w}{\left(\frac{1-p}{p}\right)^w - 1}.$$

The final solution to the recurrence is then

$$\begin{aligned} X_n &= \left\langle \frac{-\left(\frac{w}{1-2p}\right)\left(\frac{1-p}{p}\right)^n}{\left(\frac{1-p}{p}\right)^n - 1} + \left\langle \frac{\left(\frac{w}{1-2p}\right)\left(\frac{1-p}{p}\right)^n}{\left(\frac{1-p}{p}\right)^n - 1} + \frac{n}{1-2p} \right. \right. \\ &= \frac{n}{1-2p} - \left\langle \left(\frac{w}{1-2p}\right) \left[\frac{\left(\frac{1-p}{p}\right)^n - 1}{\left(\frac{1-p}{p}\right)^w - 1} \right] \right. \end{aligned}$$

Yikes! The gambler won't have any fun at all if she is thinking about this equation. Let's see if we can make it simpler in the case when $m = w - n$ is large.

Since $p < 1/2$, $(1-p)/p > 1$ and for large m ,

$$\lim_{m \rightarrow \infty} \left(\frac{w}{1-2p}\right) \left[\frac{\left(\frac{1-p}{p}\right)^n - 1}{\left(\frac{1-p}{p}\right)^w - 1} \right] = \left\langle \lim_{m \rightarrow \infty} \left(\frac{w}{1-2p}\right) \left(\frac{1-p}{p}\right)^{-m} = 0.$$

This means that as m gets large,

$$X_n \sim \frac{n}{1-2p},$$

which is much simpler. It says that if the gambler starts with $\$n$, she will expect to make about $n/(1-2p)$ bets before she goes home broke. This seems to make sense since she expects to lose

$$1 \cdot (1-p) + (-1)p = 1-2p$$

dollars on every bet and she started with n dollars.³

³Be careful, it is tempting to use such a direct and simple argument instead of all those nasty recurrences, but such an argument is not correct. There are examples where the expected duration of a process is not close to the starting point divided by the expected decrease at each step.




Figure 20.3 $n + 1$ people sitting in a circle. The B indicates the person with the broccoli—in this case, person 0.

20.3 Walking in Circles

So far, we have considered random walks on a line. Now we’ll look at a problem where the random walk is on a circle. Going from a line to a circle may not seem like such a big change, but as we have seen so often with probability, small changes can have large consequences that are often beyond the grasp of our intuition.

20.3.1 Pass the Broccoli

Suppose there are $n + 1$ people, numbered $0, 1, \dots, n$, sitting in a circle as shown in Figure 20.3. The B in Figure 20.3 indicates that person 0 has a big stalk of nutritious broccoli, which provides 250% of the US recommended daily allowance of vitamin C and is also a good source of vitamin A and iron. (Typical for a random walk problem, this game originally involved a pitcher of beer instead of a broccoli. We’re taking the high road again.)

Person 0 passes the broccoli either to the person on his left or the person on his right with equal probability. Then, that person also passes the broccoli left or right at random and so on. After a while, everyone in an arc of the circle has touched the broccoli and everyone outside that arc has not. Eventually, the arc grows until all but one person has touched the broccoli. That final person is declared the winner because they have avoided the broccoli for the longest time.

Suppose that you are allowed to position yourself anywhere in the circle. Where should you stand in order to maximize the probability that you win? You shouldn’t be person 0; you can’t win in that position. The answer is “intuitively obvious”: you should sit as far as possible from person 0, which would be position $n/2$ or $(n + 1)/2$ depending on whether n is even or odd.

20.3.2 There Is No Escape

Let’s try to verify this intuition. Suppose that you sit at position $k \neq 0$. At some point, the broccoli is going to end up in the hands of one of your neighbors. This has to happen eventually; the game can’t end until at least one of them touches it. Let’s say that person $k - 1$ gets the broccoli first. Now let’s cut the circle between yourself and your other neighbor, person $k + 1$:

$$k \quad (k - 1) \quad \dots \quad 3 \quad 2 \quad 1 \quad 0 \quad n \quad (n - 1) \quad \dots \quad (k + 1).$$

B

There are two possibilities. If the broccoli reaches you before it reaches person $k + 1$, then you lose. But if the broccoli reaches person $k + 1$ before it reaches you, then every other person has touched the broccoli and you win. So we need to compute the probability that the broccoli hops $n - 1$ people to the right before it takes 1 hop to the left. This will be the probability that you win.

But this is just the flea problem all over again. From the analysis in Section 20.1.3, we know that the probability of moving $n - 1$ steps rightward before moving one step leftward is simply $1/n$. This means that wherever you sit (aside from position 0, of course), your probability of getting the broccoli last is $1/n$.

So our intuition was completely wrong (again)! It doesn’t matter where you sit. Being close to the broccoli or far away at the start makes no difference; there is no escape—you still get the broccoli last with probability $1/n$.

Enough with the bad news: Stencil’s doomed, you go home broke from the casino, and you can’t escape the broccoli. Let’s see how to use probability to *make* some money.

MIT OpenCourseWare
<http://ocw.mit.edu>

6.042J / 18.062J Mathematics for Computer Science
Fall 2010

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.