Gambler's Ruin - Part 2

Math 394

Armed with the concept of random variables, we can inquire about an interesting question in the "gambler's ruin" (aka random walk) problem. We know that, with probability one, the game will end when both players have limited budgets, and we know what the probability of each one winning is. But how about the *time it takes to complete the game*?

Hitting Time

For simplicity, we will concentrate on the case of a "fair game", where both players have probability $\frac{1}{2}$ to win any given game. The general case can be treated similarly, but requires a little more care. Consider the problem. We have a potentially infinite sequence of independent games. Taking the point of view of player one, this can be described as a sequence of independent random variables X_i, with $P[X_i = 1] = P[X_i = -1] = \frac{1}{2}$ – i.e., player one gains a dollar for each win, and loses one for each loss. Let $X_1 + X_2 + \ldots + X_n$ be the balance between wins and losses after n games. If player one starts with x dollars (positive or negative), after n games her total amount of dollars will be

$$S_n(x) = x + \sum_{i=1}^{n} X_i$$

The game ends in T steps where T is the first integer such that $S_T = a$ or $S_T = b$, whichever happens first. Player one wins if $S_T = b$, and loses if $S_T = a$. The problem is interesting only if $a \leq x \leq b$, with $T = 0$ if $x = a$ or $x = b$.

T is (reasonably enough) called a "hitting time", as it is the time when the random walk hits one of two fixed levels.

Consider two questions, in order of difficulty:

1. What is the value of $E[T]$?
2. Can we describe the probability mass function of T?

1 Expected Hitting Time

The trick to find $E[T]$ is very similar to the one used to find the probability of win or ruin that we discussed earlier. There is a slight adjustment to make, as we are now counting how many steps we will need. Note that $E[T]$ will depend on the starting point x. As before, we consider the two possibilities for the outcome of the first game: player one either wins ($X_1 = 1$), or loses ($X_1 = -1$). Since conditional probabilities are probabilities, we can compute expectations with respect to them. We will consider the time it takes to reach one of the two boundaries, *conditional on starting from* x, and write probabilities and expectation, conditional on this, as P_x and E_x, as shorthand (this is a common notation in this kind of problems).

After the first game, we will be either in $x+1$ or in $x-1$, and, by independence, the "game starts from scratch" from this point. However, one unit of time has been used up. Thus, conditioning on the two possible outcomes of the first game,

$$E_x[T] = 1 + E_{x+1}[T]\frac{1}{2} + E_{x-1}[T]\frac{1}{2}$$

$$\frac{1}{2}\left(E_{x+1}[T] - E_x[T]\right) + \frac{1}{2}\left(E_{x-1}[T] - E_x[T]\right) = -1$$

You will recognize that the same calculation will work at any time step. Thus, if we define

$$u(x) = E_x[T]$$

we have the difference equation

$$\frac{1}{2}\left[u(x+1) - u(x)\right] + \frac{1}{2}\left[u(x-1) - u(x)\right] = -1, \quad u(a) = u(b) = 0 \tag{1}$$

We could try to guess what a solution to this equation could look like, but to help our intuition, let's take a limit similar to the one we took in the original gambler's ruin problem, which led us to a differential equation. That is, let's change the size of the step to a small number, say h.

Now, if the steps are shorter, we need to walk faster. That is the time step has to become shorter too – no longer 1. If we choose a time step also equal to h we don't get very far:

$$\frac{1}{2}\left[u(x+h) - u(x)\right] + \frac{1}{2}\left[u(x-h) - u(x)\right] = -h$$

$$\frac{1}{2}\left[\frac{u(x+h) - u(x)}{h}\right] + \frac{1}{2}\left[\frac{u(x-h) - u(x)}{h}\right] = -1$$

and in the limit $h \to 0$, if we assume u to be differentiable, we end up with $0 = 1$, meaning the equation does not have a limit.

Let's then try h^2. This time we are on to something:

$$\frac{1}{2}\left[u(x+h)-u(x)\right]+\frac{1}{2}\left[u(x-h)-u(x)\right]=-h^2$$

$$\frac{1}{2}\left[\frac{u(x+h)-u(x)}{h^2}\right]+\frac{1}{2}\left[\frac{u(x-h)-u(x)}{h^2}\right]=-1$$

and with $u(x+h)-u(x)\approx u'(x)\cdot h+\frac{1}{2}u''(x)\cdot h^2$, $u(x-h)-u(x)\approx -u'(x)\cdot h+\frac{1}{2}u''(x)\cdot h^2$, our equation has a proper limit:

$$\frac{1}{2}u''(x)=-1, \quad u(a)=u(b)=0 \tag{2}$$

(This is a one-dimensional version of the so-called *Dirichlet problem for a Poisson equation*). Solving $u''=-2$ yields, as a general solution $u(x)=-x^2+Ax+B$, where A and B are determined by the boundary conditions, that is that a and b be the roots of the polynomial, i.e.

$$-x^2+Ax+B=-\left(x-a\right)\left(x-b\right)$$

that is

$$A=(a+b), \quad B=-ab$$

It is not hard to check that we will get the same result if we plug a tentative solution of the form $-x^2+Ax+B$ into equation (1), or, equivalently, we check that $-x^2+(a+b)\,x-ab$ solves (1).

Consider now the limit case as $a\to-\infty$. This would result in the time to reach a barrier of height b with no lower bound. Since the solution to our equation (2) can only be quadratic, we see that there is no solution in this case. If we look at the solution

$$u(x)=-x^2+(a+b)\,x-ab$$

for fixed x, it will look like $a(x-b)$, which, as $a\to\infty$, diverges to ∞. In other words, even though the barrier $x=b$ will be reached with probability 1 (in fact, choosing $x=b-1$, it will be reached in just one step with probability $\frac{1}{2}$), *on average it will take an infinite number of steps to reach it.*

Note: Feller, in his fundamental work on the theory of probability has a tongue-in-cheek application of this result. Suppose you are in line at a freeway toll station, and look at the cars in the line besides you. If the two lines move at random, each time line 1 moving by one step with probability $\frac{1}{2}$, if you target the car in the neighboring line one step ahead of you, your position with respect to it will be a random walk just like the one we studied. We just proved that, on average, it will take you an infinite amount of time to catch up! Which, as Feller notes, proves, as a theorem, your deep conviction that "I always end up in the slowest lane". Of course, what is really maddening is that people in the neighboring line will reach exactly the same conclusion.

2 Distribution of the Hitting Time

Let's be more ambitious, and try to find an equation for the *distribution* of the hitting time. That is, for example,

$$P_x[T > t] \equiv v(x,t) \tag{3}$$

The logic is similar to the other random walk problems we have seen. Conditioning on the outcome of the next game, we will have (we have used up one time step)

$$P_x[T > t] = \frac{1}{2}P_{x-1}[T > t-1] + \frac{1}{2}P_{x+1}[T > t-1]$$

With our notation (3), this is

$$v(x,t) = \frac{1}{2}[v(x+1,t-1) + v(x-1,t-1)] \tag{4}$$

This is a much less simple equation than our previous ones, so we move to our limit for small space and time steps. We already know that for a limit to make sense, the time steps have to be of the order of the square of the space steps. Thus our scaled random walk will go from x to $x \pm h$ in time h^2:

$$v(x,t) = \frac{1}{2}\left[v\left(x+h,t-h^2\right) + v\left(x-h,t-h^2\right)\right]$$

To find a limit we will subtract $2v\left(x,t-h^2\right)$ from both sides, and divide both sides by h^2, as this will cause familiar differences to appear:

$$\frac{v(x,t) - v\left(x,t-h^2\right)}{h^2} = \frac{1}{2}\left[\frac{v\left(x+h,t-h^2\right) + v\left(x-h,t-h^2\right) - 2v\left(x,t-h^2\right)}{h^2}\right]$$

We have, as usual, to assume that our function v will behave nicely as we let $h \to 0$, that is assume that it has a continuous derivative with respect to t, and two continuous derivatives with respect to x. This can be justified after the fact, if it turns out that the resulting equation has a nice solution (and, in fact, it turns out to have an exceptionally nice solution).

If we now let $h \to 0$, the left hand side converges to $\frac{\partial v(x,t)}{\partial t}$, and the right hand side converges as well, since the terms in h in the numerator cancel, and we end up with a second derivative:

$$\frac{\partial v(x,t)}{\partial t} = \frac{1}{2}\frac{\partial^2 v(x,t)}{\partial x^2} \tag{5}$$

Equation (5) is called the *Heat Equation*, as it has originally been introduced to describe heat conduction. Together with the *boundary conditions* $v(a) = v(b) = 0$, and an initial condition $v(x,0) = v_0(x)$ which, surprisingly, need not be "nice" at all, it can be shown that this equation has a unique solution, which is infinitely differentiable (regardless of how rough the initial condition is!), but the solution will, in general, not be expressible in closed form, but rather as a series of functions (in fact, a *Fourier Series* is the natural choice here).