Notes on the Gamber's Ruin Problem

September 30, 2018

Abstract

These are some notes on the Gambler's Ruin Problem.

0.1 The Gambler's Ruin Problem

Consider a game between two players. The game consists of a sequence of independent, identical trials in each of which there is a probability p that player A wins, and a probability of $1-p$ that player B wins. We assume $0 < p < 1$ to avoid trivialities. For example, in each trial, two standard dice could be tossed, and player A would win if the sum of the resulting numbers was 5 or 7, while player B would win if the sum was anything else. In this case $p = \frac{10}{36}$.

Now suppose players A and B have have n_0 and m_0 dollars, respectively, to wager. At each trial, the loser pays the winner one dollar. The game ends when one player is broke, so that when the game ends, one player ends up with $n_0 + m_0$ dollars, and the other with nothing. What is the probability that player A ends up with all of the money?

Define $T = m_0 + n_0$; T denote the total fortune of the two players. What is the probability that player A wins, as a function of n_0, T and p?

0.2 The probability space

The first step towards solving this problem is to write down a probability space (S, P) within which it can be solved. This takes us into new territory. Let S be the set of all infinite sequences $\{x_j\}_{j\in\mathbb{N}}$ where each x_j is either 1 or -1. Let $x_j = 1$ correspond to player A winning the jth trial.

The set S is uncountable. This is because it has the same cardinality as the set of all infinite sequences of 0's and 1's, and each such sequence that has only finitely many entries that are 1 can be viewed as specifying the canonical binary expansion of a real number in the interval $[0, 1)$, and there are uncountably many of these. The problem of defining a probability measure P on an uncountable set requires measure theory for its full solution, but we can avoid this by considering, for each N, a special class of events in S, and this is very natural to do, and useful for other purposes.

0.1 DEFINITION. A subset E of S is is *measurable at trial* N in case the question as to whether or not $\omega = \{x_j\}_{j\in N}$ does or does not belong to E can be answered knowing only $(x_1, \dots, x_N)$, and not anything about the rest of the infinite sequence specifying ω. That is, membership in the set E depends only on the outcomes of the first N trials. The set of all such events is denoted $\mathcal{F}_N$.

For example, consider the event E that player B wins the first three trials. Then

$$\omega = \{x_j\}_{j\in N} \in E \qquad \iff \qquad (x_1, x_2, x_3) = (-1, -1, -1)\ .$$

Therefore, $E \in \mathcal{F}_3$. Note also that $E \in \mathcal{F}_N$ for all $N \geq 3$, but E does not belong to $\mathcal{F}_1$ or $\mathcal{F}_2$. A little reflection will show that for any event $E \in \mathcal{F}_M$ for some M, $E \in \mathcal{F}_N$ for $N \geq M$.

The point of the definition is that it is easy to assign probabilities to events that belong to $\mathcal{F}_N$ for some N. As an example, consider the event E that player B wins the first three trials. By the independence of the trials, and the fact that it does not matter, as far as membership in E is concerned, what happens after the first three trials, we must have that

$$P(E) = (1-p)^3\ .$$

More generally, for $N \in \mathbb{N}$, let S_N denote the set of vectors $\widehat{\omega} = (x_1, \dots, x_N)$ where each x_j is either 1 or -1, interpreted as before. There are 2^N outcomes in S_N, and by independence, if $\widehat{\omega}$ has k positive terms, the probability of $\widehat{\omega}$, which we shall denote by $\widehat{P}(\widehat{\omega})$, is given by $\widehat{P}(\widehat{\omega}) = p^k(1-p)^{N-k}$. For any event $\widehat{E} \subset S_N$, we then have

$$\widehat{P}(\widehat{E}) = \sum_{\widehat{\omega}\in E} \widehat{P}(\widehat{\omega})\ . \tag{0.1}$$

Then $(S_N, \widehat{P})$ is a well defined finite probability space.

0.2 EXAMPLE. *Consider once more the event E that player B wins the first three trials, which belongs to $\mathcal{F}_3$ and $\mathcal{F}_4$, and for that matter to $\mathcal{F}_N$ for all $N \geq 3$, since once you know the result of the first 3 trials, you know whether the event occurred or not. Hence, knowledge of the results of the first 4 or more trails is overkill.*

If we apply formula (0.1) in S_3, we have

$$\widehat{E} = \{(-1, -1, -1)\} \quad \text{and} \quad P(E) = \widehat{P}(\widehat{E}) = (1-p)^3\ .$$

If we apply formula (0.1) in S_4, we have

$$\widehat{E} = \{(-1, -1, -1, 1), (-1, -1, -1, -1)\} \quad \text{and} \quad P(E) = \widehat{P}(\widehat{E}) = (1-p)^3 p + (1-p)^4 = (1-p)^3\ .$$

A bit of reflection on the example will how that for *any* event $E \in \mathcal{F}_M$ for some M, then one can compute $P(E)$ by applying (0.2) in S_N for any $N \geq M$, and the result is always the same.

Now given $E \in \mathcal{F}_N$, define $\widehat{E}$ to be the subset of S_N such that

$$\omega = \{x_j\}_{j\in N} \in E \qquad \iff \qquad \widehat{\omega} = (x_1, \dots, x_N) \in \widehat{E}\ .$$

That is, $\widehat{E}$ is precisely the set of outcomes of the first N trials that qualify for membership in E, which makes sense because membership in E *depends only on the outcomes of the first* N *trials* by the very definition of $\mathcal{F}_N$. Now define our probability measure P on the events E that belong to $\mathcal{F}_N$ for some $N \in \mathbb{N}$ by

$$P(E) = \widehat{P}(\widehat{E}) \ . \tag{0.2}$$

There is no N on the right side for the following reason: If $E \in \mathcal{F}_N$ for some N, then there is a least N_0 such that $E \in \mathcal{F}_{M_0}$, and then $E \in \mathcal{F}_M$ for all $M \geq N_0$. But by what we have explained in Example 0.2 and the remark following it, $\widehat{P}(\widehat{E})$ may be computed in $(S_M, \widehat{P})$ for any $M \geq N_0$, and the result is always the same.

0.3 EXAMPLE. *Let U_N be the event that the game is undecided after N trials. Note that $\omega = \{x_j\}_{j\in\mathbb{N}}$ belongs to U_N if and only if*

$$-n_0 < \sum_{j-1}^{\ell} x_j < m_0 \tag{0.3}$$

for all $\ell \leq N$, since this is precisely the condition that neither player has cleaned the other player out by the Nth trial. Clearly $U_N \in \mathcal{F}_N$, since membership in U is determined by the first N trials, and $\widehat{U_N} \subset S_N$ is given by the set of vectors $(x_1, \dots, x_N)$ that satisfy (0.3). Thus, the probability of U_N, $P(U_N)$, is given by $P(U_N) = \widehat{P}(\widehat{U_N})$, and this probability is given by a finite sum.

We now claim that $\lim_{N\to\infty} P(U_N) = 0$. In fact, it converges to zero quite fast, and we can be fairly extravagant with our upper bounds. Take N much larger that m_0 and n_0. Let E_1 be the event that the first $m_0 + n_0$ trials all result in -1. Let E_2 be the event that the second $m_0 + n_0$ trials all result in -1, and so forth. Let M be the largest integer such that $M(m_0+n_0) \leq N$. Then the events $\{E_1, \dots, E_M\}$ are all in $\mathcal{F}_N$, and the events $\{\widehat{E_1}, \dots, \widehat{E_M}\}$ are mutually independent in $(S_N, \widehat{P})$. Also, again by the independence of the trials, $\widehat{P}(\widehat{E_j}) = (1-p)^{m_0}$ for each $j = 1, \dots, M$. Thus, $P(E_j) = (1-p)^{m_0+n_0}$ for each $j = 1, \dots, M$.

Clearly, if any of the events E_j happen, the game is decided. Therefore

$$\bigcup_{j=1}^{M} E_j \subset U_N^c \ ,$$

and hence

$$U_N \subset \left(\bigcup_{j=1}^{M} E_j\right)^c = \bigcap_{j=1}^{M} E_j^c \ .$$

It follows that

$$P(U_N) \leq (1-(1-p)^{m_0+n_0})^M \ .$$

and since $M \to \infty$ and $\mathbb{N} \to \infty$, this proves that $P(U_N)$ converges to zero.

Now let us consider the event that is most relevant to the Gambler's Ruin problem, namely the event W that player A wins the game. Note that $\omega = \{x_j\}_{j\in\mathbb{N}}$ belongs to W if and only if for some $k \in \mathbb{N}$,

$$\sum_{j-1}^{k+1} x_j = m_0 \ , \tag{0.4}$$

and

$$-n_0 < \sum_{j-1}^{\ell} x_j < m_0 \tag{0.5}$$

for all $\ell < k$. This is because (0.5) says that no player has cleaned the other player out up through the kth trial, so nobody has already one before the kth trial, but then by (0.4) player A cleans out player B on the kth trial. Since k can be arbitrarily large, W *does not* belong to $\mathcal{F}_N$ for any $N \in \mathbb{N}$.

All is not lost. As we have seen the probability that the game is not over by the Nth trial goes to zero, exponentially fast even, as N increases. Define W_N to be the event that player A has won the game by the Nth trial, Evidently $W_N \in \mathcal{F}_N$, and so this event has a well-defined probability. Moreover it is clear that for all $M < N$, $W_M \subset W_N$, because if player A has won by the Mth trial they have won by the Nth, but even if they haven't ,they might still win in the next $N - M$ trials. Therefore, $P(W_N)$ is an increasing function of N. Furthermore, since whenever player A does win, they must have won by the Nth trial for some N, and hence

$$W = \bigcup_{N \in \mathbb{N}} W_N \ .$$

That is, we have

$$W_1 \subset W_2 \subset \cdots \mathrm{W}_N \subset W_{N+1} \subset \cdots \subset W \ ,$$

and the union of all of the sets on the left is not only contained in W, it equals W. In this circumstance, we define

$$P(W) = \lim_{N \to \infty} P(W_N) \ ,$$

where the limit exists because bounded monotone sequences always converge.

Now we can give a proper definition of the quantity of interest: For any $m, n, N \in \mathbb{N}$, let W_N be the event that player A wins by the Nth trial given that the initial fortune of player A is $n_0 = n$, and the initial fortune of plater B is $m_0 = m$. Let T denote $m_0 + n_0 = m + n$. Then for this event W_N, define

$$P_{n,T,N} = P(W_N) \ ,$$

and define

$$P_{n,T} = \lim_{N \to \infty} P_{n,T,N} \ .$$

where both quantities are properly defined by what has been explained above.

0.3 Recursion for $P_{n,T}$

We are now ready to tackle the Gambler's Ruin problem and answer the question posed at the beginning of tis section of notes. Let E be the event that player A wins the first trial. Suppose $0 < n < T$. Then, with W_N defined exactly as in paragraph above,

$$W_N = (W_N \cap E) \cup (W_N \cap E^c)$$

and since the events on the right are mutually exclusive,

$$P(W_N) = P(W_N \cap E) + P(W_N \cap E^c)\ .$$

Notice that the probabilities are well-defined since all of the events in question belong to $\mathcal{F}_N$. Now to take advantage of the independence, we write this in terms of conditional probabilities:

$$P(W_N) = P(E)P(W_N|E) + P(E^c)P(W_N|E^c)\ .$$

Given the event E, after the first trial, player A has $n+1$ dollars, and player B has $M-1$. T is unchanged, and player A has $N-1$ more trials in which to win (if the game does not end earlier). Since the trials are independent, the game starts afresh from this new starting point, and

$$P(W_N|E) = P_{n+1,T,N-1}\ .$$

The same reasoning yields $P(W_N|E^c) = P_{n-1,T,N-1}$, and of course $P(E) = p$ and $P(E^c) = 1-p$. Altogether,

$$P_{n,T,N} = P(W_N) = pP_{n+1,T,N-1} + (1-p)P_{n-1,T,N-1}\ .$$

Now taking $N \to \infty$, we obtain

$$P_{n,T} = pP_{n+1,T} + (1-p)P_{n-1,T}\ . \tag{0.6}$$

So far, we have supposed that $0 < n < T$. But if $n = 0$, player A cannot win, so $P_{0,T} = 0$, while if $n = T$, player A has already won, so $P_{0,T} = 1$.

0.4 LEMMA. *For any $T \geq 2$, $p \neq 1/2$, a vector $(f_1, \dots f_T)$ satisfies*

$$f_n = pf_{n-1} + (1-p)f_{n+1} \tag{0.7}$$

for all $0 < n < T$ if and only if, for some constants α and β

$$f_n = \alpha + \beta\left(\frac{1-p}{p}\right)^n \tag{0.8}$$

for all $0 \leq n \leq T$. Moreover, the values of α and β are uniquely determined.

Proof. Let us seek a solution of (0.7) of the form $f_n = a^n$, with a to be chosen. Then (0.7) becomes

$$a^n = pa^{n+1} + (1-p)a^{n-1} \ .$$

Dividing by a^{n-1}, we obtain the quadratic equation

$$a^2 - \frac{1}{p}a + \frac{1-p}{p} = 0 \ .$$

The two roots are $a = 1$ and $a = \frac{1-p}{p}$. As long as $p \neq 1/2$, this gives us two distinct solutions, namely $f_n = 1$ for all n, and $f_n = \left(\frac{1-p}{p}\right)^n$ for all n. Since the equation is linear, every linear combination of solutions is a solution, and hence (0.8) does indeed specify a solution.

It remains to show that every solution has this form for exactly one choice of α and β. Let $(\tilde{f}_0, \dots, \tilde{f}_T)$ be any such solution. Let $(f_0, \dots, f_T)$ be given by (0.8). We seek to choose α and β so that $f_0 = \tilde{f}_0$ and $f_1 = \tilde{f}_1$. This will imply that $f_n = \tilde{f}_n$ for all n, since the recursion equation can be written as

$$f_{n+2} = \frac{1}{p}f_{n+1} - \frac{1-p}{p}f_n$$

for all $0 < n < T-2$. It is clear from this that once the first two terms are known, all the rest are determined. Therefore, we try to solve $f_0 = \tilde{f}_0$ and $f_1 = \tilde{f}_1$, which gives us

$$\begin{aligned} \alpha + \beta &= \tilde{f}_0 \\ \alpha + \beta\frac{1-p}{p} &= \tilde{f}_1 \ . \end{aligned}$$

Subtracting the first equation from the second, $\beta\frac{1-2p}{p} = \tilde{f}_1 - \tilde{f}_0$, and thus the unique solution is

$$\beta = \frac{p}{1-2p}\tilde{f}_1 - \tilde{f}_0 \quad \text{and} \quad \alpha = \beta - \tilde{f}_0 \ .$$

$\square$

To apply this, we see that the vector $(P_{0,T}, P_{1,T}, \dots, P_{T,t})$ is a solution of (0.7), and hence it must be given by (0.8) for some uniquely determined α and β. That is,

$$P_{n,T} = \alpha + \beta\left(\frac{1-p}{p}\right)^n$$

for all $0 \leq n \leq T$. Since we also know that $P_{0.0} = 0$, we see that $\alpha + \beta = 0$, so we eliminate α to find

$$P_{n,T} = \beta\left(-1 + \left(\frac{1-p}{p}\right)^n\right)$$

Next, we also know that $P_{T,T} = 1$, and therefore

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

which specifies β. Finally we have proved:

0.5 THEOREM. *With $P_{n,T}$ defined as above, for $0 \le n \le T$,*

$$P_{n,T} = \frac{\left(\frac{1-p}{p}\right)^n - 1}{\left(\frac{1-p}{p}\right)^T - 1} \ . \tag{0.9}$$

0.6 EXAMPLE. *Consider the problem raised at the beginning. Player A has a probability of $p = \frac{10}{36}$ of winning each trial, but starts with a fortune of $n = 100$ dollars. Player B starts with a fortune of $m = 10$ dollars. What is the probability that A wins?*

In this case, $n = 100$, $T = 110$ and $p = \frac{10}{36}$, and $(1-p)/p = \frac{13}{5}$, so the answer is

$$P_{100,110} = \frac{(\frac{13}{5})^{100} - 1}{(\frac{13}{5})^{110} - 1} = 0.00070838037.... \ .$$

The chance that A wins, even with their much deeper pockets, is very, very small.

There is more to learn from the last example. Notice that $(\frac{13}{5})^{100}$ is really large and $(\frac{13}{5})^{110}$ is even larger, and so the -1 terms hardly matter. In fact, for $b > a > 1$, and $t \le 1$,

$$\frac{\mathrm{d}}{\mathrm{d}t}\frac{a-t}{b-t} = \frac{a-b}{(b-t)^2} \le 0 \ .$$

Thus,

$$\frac{a-1}{b-1} \le \frac{a}{b} \ .$$

It follows that:

0.7 LEMMA. *whenever $p < 1/2$,*

$$P_{n,T} \le \left(\frac{p}{1-p}\right)^m \tag{0.10}$$

independent of $n \ge 1$.

Proof. By what we have noted just above,

$$P_{n,T} = \frac{\left(\frac{1-p}{p}\right)^n - 1}{\left(\frac{1-p}{p}\right)^T - 1} \le \frac{\left(\frac{1-p}{p}\right)^n}{\left(\frac{1-p}{p}\right)^T} = \left(\frac{p}{1-p}\right)^m \ .$$

□

Therefore, if you could convince the richest person in the world, Jeff Bezos at the time of writing, to play this game, and you had just 10 dollars to stake again all of their fortune, you would have a probability of no less than $1 - (\frac{5}{13})^{10} = 0.999929...$ of winning their entire fortune. Of course if each trial takes only one second, the minimum amount of time it will take you to win your first billion is 10^9 seconds, which is about 31.7 years. And that is if you win *all* of the first billion trials, which is *not* going to happen.

What we have seen so far, is that in such a game, the important thing is to stake out a position where the odds on each trial are in your favor. Then even if you start with a modest fortune, you have a good chance of cleaning out an arbitrarily rich opponent. However, will you live long enough to do this? How long will it take you to make your first billion in such a wager? We shall return to these questions when we have more theoretical tools to handle them. For now, we close with one more example:

0.8 EXAMPLE. *Suppose* $p = 0.49$ *so the odd are only slightly in favor of player B. Suppose the initial fortune of player B is* $m = 1$ *dollar, and the initial fortune of player A is* $n = 100$ *dollars. Using the formulas, we find that the probability that A wins is only* 0.96008224.... *There is an almost* 4% *chance that player B cleans player A out. And with* $m = 18$ *and* $n = 100$*, the chance of B cleaning A out is over half.* Even a slight tilt in the odds makes a big difference.

0.4 The case $p = 1/2$

When $p = 1/2$, we still have the recursion relation (0.6), which becosmes

$$P_{n,T} = \tfrac{1}{2}P_{n+1,T} + \tfrac{1}{2}P_{n-1,T}\ . \tag{0.11}$$

As before, we seek all vectors $(f_0, \ldots, f_T)$ satsifying

$$f_n = \frac{1}{2}(f_{n+1} + f_{n-1}) \tag{0.12}$$

for all $0 < n < T$.

We always have the constant solution $f_n = 1$ for all n; this is a solution for every choice of p. You might be able to guess the second solution, especially if you recognize the equation as a discrete approximation to the differntial equation $f''(x) = 0$. However, let's derive the second solution by a limiting procedure. This will make it clear how the case $p = 1/2$ meshes with the cases $p \neq 1/2$.

For $p \neq 1/2$, we found the general solution

$$f_n = \alpha + \beta\left(\frac{1-p}{p}\right)^n ,$$

for constants α and β. If we take the limit $p \to 1/2$ for fixed α and β, we just get the triavial constant solution. To get a linearly indepndent solution, take $\alpha = -\frac{1}{p-1/2}$ and $\beta = \frac{1}{p-1/2}$. This gives us the solution

$$f_n = \frac{1}{p-\frac{1}{2}}\left(\left(\frac{1-p}{p}\right)^n - 1\right)\ .$$

For each $p \neq 1/2$, this solves the equation (0.7) by Lemma 0.4. Now lets take the limit $p \to 1/2$. This ammounts to taking a derivative: Define

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

Then $\varphi(1/2) = 1$, so that the limit

$$\lim_{p \to \frac{1}{2}} \frac{1}{p - \frac{1}{2}} \left(\left(\frac{1-p}{p} \right)^n - 1 \right) = \varphi'(\tfrac{1}{2}) = -4n \ .$$

We can ignore the multiple -4, and thus our second solution is $f_n = n$, and the general solution of (0.12) is

$$f_n = \alpha + \beta n \ .$$

If we require $f_0 = 0$, then $\alpha = 0$, and if we require $f_T = 1$, then $\beta = \frac{1}{T}$. Therefore, for $p = 1/2$,

$$P_{n,T} = \frac{n}{T} \ .$$

For $p = 1/2$, the chance that A wins *is proportional to their initial fortume.*