Symbol-Crunching With the Gambler's Ruin Problem

Doron ZEILBERGER[1]

Dedicated to my brother Gil Zeilberger

Prelude:
George may be a Better Backgammon Player than Martha, BUT He is not That Much Better!

My good friend George and his wife Martha play lots and lots of Backgammon. When I inquired who was better, George replied proudly: "Of course I am so much better, so far I got 250 *stars*, while Martha only got 5".

When I asked what is a 'star'?, they replied that they play too many games to keep exact score, and instead only keep track of the *difference*. As soon as one of the players is 20 points ahead of the other, he or she wins a star, and they start all over.

George concluded that since the 'star-score' is 250 : 5, it follows that he is $250/5 = 50$ times better player than Martha.

But is George really *that* much better than Martha? I soon realized that this is essentially the good-old *Gambler's Ruin* problem, that goes back to Christiaan Huygens in his 1657 treatise on probability.

George and Martha do not use the doubling cube, and assume, for the sake of simplicity, that there are no gammons , i.e. every game is either a simple win for George or a simple win for Martha. In that case, their game is equivalent to two players that start with 20 dollars each and at each time-unit, give each other one dollar with probability p and $q := 1 - p$ respectively, until one of them goes broke.

According to the classical formula discovered independently by James Bernoulli and Abraham De Moivre, if currently the first player has a dollars and the second player has b dollars, then the probability of the first player winding up a winner (with $a + b$ dollars in his or her pocket) is:

$$\begin{cases} \frac{1-(q/p)^a}{1-(q/p)^{a+b}}, & \text{if } p \neq \frac{1}{2}; \\ \frac{a}{a+b}, & \text{if } p = \frac{1}{2}. \end{cases}$$

Since here $a = b = 20$, the probability of George winning a star is

$$\frac{1-(q/p)^{20}}{1-(q/p)^{40}} \quad .$$

[1] Department of Mathematics, Rutgers University (New Brunswick), Hill Center-Busch Campus, 110 Frelinghuysen Rd., Piscataway, NJ 08854-8019, USA. `zeilberg at math dot rutgers dot edu` ,
`http://www.math.rutgers.edu/~zeilberg` . Sept. 15, 2006. Accompanied by Maple package `RUIN` downloadable from Zeilberger's website. Supported in part by the NSF.

For example, if $p = .55$ then the probability of George winning is .9822, which is consistent with the 250 : 5 star-score. So Martha is a very decent Backgammon player, she wins %45 of the games!

De Moivre also proved that for the random variable 'duration of the game', let's call it D, the expectation is given by

$$E[D] = \begin{cases} \frac{b}{p-q} - \frac{a+b}{p-q}\frac{1-(p/q)^b}{1-(p/q)^{a+b}}, & \text{if } p \neq \frac{1}{2}; \\ ab, & \text{if } p = \frac{1}{2}. \end{cases}$$

So, on the average, it takes George and Martha roughly $20/(.55 - .45) = 200$ games to complete a star.

Interlude: Sermon

The Founding Fathers of probability theory were much smarter than me, but I have *one* thing that they didn't have, I own a computer! More important, my computer has a computer algebra system (Maple in my case). Can I do more?

Of course I can! Or rather my computer can. Probability Theory, (and mathematics in general,) evolved the way they did because in those days people didn't have computers, and some questions they didn't even dare ask, since they realized that it would be hopeless to find exact answers. So they had to resort to limiting theorems, asymptotics and dubious approximations. In particular, they largely abandoned discrete random variables and worked directly with continuous ones, and they often gave up on "=" in favor of "$\sim$" and "$\leq$".

I believe that it would be a very rewarding, and potentially useful, intellectual exercise to start probability theory *ab initio*. Let's forget about the 300 years of computer-less human development, and start all over, taking full advantage of computers.

Now I don't mean, as Stephen Wolfram advocates, to dump *equations* altogether, and just run computer simulation à la Monte Carlo. After all, it is a bit boring to just be passive watchers. Let's keep **formulas** and **exact answers** (whenever possible), but let the computer derive them!

General Gambler's Ruin

Going back to the Gambler's Ruin problem, what questions can we ask that De Moivre would never be able to answer?

Sure enough he found the *expectation* (alias *first* moment), of D (the 'duration of the game' r.v.), but what about its *variance*?, its *skewness*?, its *kurtosis*?, its *tenth moment*?

Of course, we can naively use *simulation*, but we can't get exact answers that way, not even numerically. To get *exact* numerical answers (for each specific (numerical) $A := a + b$), and any specific p, we can set-up the Markov Chain with the two absorbing states 'George is broke' and 'Martha is broke', and then set-up the (numerical) transition matrix, solve the system of linear equations etc. etc.

But we want *symbolic* answers, phrased in terms of a, b and p, and for the general die, in terms of the probability distribution for the die's faces. In other words, we would like to have **general** formulas valid for *arbitrary* a and b, that we can just simply plug-in without doing it every time from scratch.

These problems are way beyond the scope of mere humans, but they can be easily handled by machines. At this time of writing, however, we still need a human to *design* the algorithm. This will be done in the next section.

The Algorithm

For the sake of exposition, let's only do the fair coin case, that is equivalent to a random walker starting when $t = 0$ at $x = n$, and at every time-step, if he is at location x at time t, then, at time $t+1$, he is going to be in location $x+1$ with probability $1/2$ and in location $x-1$ with probability $1/2$. The game terminates as soon as he visits $x = 0$ or $x = A$.

Let $P(t, n)$ be the probability of finishing exactly after t steps, if it starts at $x = n$.

Obviously, $P(t, n)$ satisfies, for $t > 0$ and $0 < n < A$, the linear partial recurrence

$$P(t,n) = \frac{1}{2}P(t-1,n-1) + \frac{1}{2}P(t-1,n+1) \quad .$$

Indeed, if our walker is at position n today, then tomorrow he would be at position $n+1$, or position $n-1$, each with probability $1/2$, and in order to finish t days from today, he would have to finish $t-1$ days from tomorrow.

We also need the *initial condition*

$$P(0,n) = \begin{cases} 1, & \text{if } n = 0 \text{ or } n = A \ ; \\ 0, & \text{if } 0 < n < A \ , \end{cases}$$

and the *boundary conditions*

$$P(t,0) = \begin{cases} 1, & \text{if } t = 0 \ ; \\ 0, & \text{if } t > 0 \ , \end{cases}$$

$$P(t,A) = \begin{cases} 1, & \text{if } t = 0 \ ; \\ 0, & \text{if } t > 0 \ . \end{cases}$$

Let's first try and compute $F(n)$, the probability that the game will eventually end. By definition,

$$F(n) = \sum_{t=0}^{\infty} P(t,n) \quad .$$

We have

$$F(n) = P(0,n) \, + \sum_{t=1}^{\infty} P(t,n) = P(0,n) \, + \sum_{t=1}^{\infty} \left(\frac{1}{2}P(t-1,n-1) + \frac{1}{2}P(t-1,n+1) \right)$$

$$= P(0,n) \ + \ \frac{1}{2}\sum_{t=1}^{\infty} P(t-1,n-1) + \frac{1}{2}\sum_{t=1}^{\infty} P(t-1,n+1)$$

$$= P(0,n) \ + \ \frac{1}{2}\sum_{t=0}^{\infty} P(t,n-1) + \frac{1}{2}\sum_{t=0}^{\infty} P(t,n+1) = P(0,n) \ + \ \frac{1}{2}F(n-1) \ + \ \frac{1}{2}F(n+1) \quad .$$

We thus have the *linear recurrence*

$$F(n) = \frac{1}{2}F(n-1) \ + \ \frac{1}{2}F(n+1) \quad , \quad 0 < n < A \quad ,$$

with the boundary conditions $F(0) = 1, F(A) = 1$. Its unique solution is $F(n) \equiv 1$. So we have just proved that with probability 1 the game will eventually end.

Now, let's try and compute the r-th *binomial moment*, let's call it $f_r(n)$, of our "duration of the game" random variable D, for $r \geq 1$. By definition,

$$f_r(n) = \sum_{t=0}^{\infty} \binom{t}{r} P(t,n) \quad .$$

We have

$$f_r(n) = \binom{0}{r} P(0,n) \ + \sum_{t=1}^{\infty} \binom{t}{r} P(t,n) = 0 \ + \sum_{t=1}^{\infty} \binom{t}{r} \left(\frac{1}{2}P(t-1,n-1) + \frac{1}{2}P(t-1,n+1)\right)$$

$$= \frac{1}{2}\sum_{t=1}^{\infty} \left(\binom{t-1}{r} + \binom{t-1}{r-1}\right) P(t-1,n-1) + \frac{1}{2}\sum_{t=1}^{\infty} \left(\binom{t-1}{r} + \binom{t-1}{r-1}\right) P(t-1,n+1)$$

$$= \frac{1}{2}\sum_{t=0}^{\infty} \binom{t}{r} P(t,n-1) + \frac{1}{2}\sum_{t=0}^{\infty} \binom{t}{r-1} P(t,n-1) + \frac{1}{2}\sum_{t=0}^{\infty} \binom{t}{r} P(t,n+1) + \frac{1}{2}\sum_{t=0}^{\infty} \binom{t}{r-1} P(t,n+1)$$

$$= \frac{1}{2}f_r(n-1) \ + \frac{1}{2}f_{r-1}(n-1) \ + \ \frac{1}{2}f_r(n+1) + \ \frac{1}{2}f_{r-1}(n+1)$$

We thus have the linear (*inhomogeneous*) recurrence

$$f_r(n) - \frac{1}{2}f_r(n-1) - \frac{1}{2}f_r(n+1) = \frac{1}{2}f_{r-1}(n-1) + \frac{1}{2}f_{r-1}(n+1) \quad , \quad 0 < n < A \quad ,$$

with the boundary conditions $f_r(0) = 0, f_r(A) = 0$.

For $r = 1$ (the expected duration of the game), we can still do it by hand, and get $E[D] = f_1(n) = n(A-n)$, which is the unique solution of the recurrence

$$X(n) - \frac{1}{2}X(n-1) - \frac{1}{2}X(n+1) = 1 \quad , 0 < n < A \quad .$$

subject to $X(0) = X(A) = 0$. But already for $r = 2$, it is getting a bit tedious, and for higher moments we really need computer algebra.

Enter the computer

Since the roots of the indicial equation of the recurrence are $\{1, 1\}$, it follows that $f_r(n)$ is a polynomial in n, of degree two more of that of $f_{r-1}(n)$. Once you know $f_{r-1}(n)$ we (or rather our computer) finds a **particular solution** obtained by setting an ansatz of a polynomial of the right degree $(2r)$, with *undetermined* (generic) coefficients, plugging it into the linear recurrence, and solving the system of equations obtained by equating all the coefficients of the powers of n, and plugging back. Then, as in Discrete Math 101, add it to the General Solution of the homogeneous version, $c_0 + c_1 n$, and plug-in $n = 0$ and $n = A$, set them both equal to 0, solve for c_1 and c_2, and plug them back.

But the **binomial moments** are only stepping stones for the actual moments. We can now ask our beloved electronic servant to find the coefficients $b_{r,s}$ such that

$$n^r = \sum_{s=0}^{r} b_{r,s} \binom{n}{s}$$

[these coefficients are Stirling numbers of one kind or another (I forgot which), but it does not matter, Maple can find them **ab initio**].

Using the $b_{r,s}$, our computer program finds

$$E[D^r] = \sum_{s=0}^{r} b_{r,s} f_r(n) \quad ,$$

and finally, the true objects of desire, the *moments about the mean*, are computed the usual way

$$m_r(D)(n) = E[(D - m_1)^r](n) = \sum_{s=0}^{r} \binom{r}{s} (-m_1)^{s-r} E[D^s] \quad ,$$

once again completely internally and seamlessly.

How long would it take to win?

If you are only interested in your probability of quitting a winner and how long you would expect that it should take and the variance and higher moments *relative* to the event 'I am ultimately going to be the winner', the previous analysis goes verbatim, with only one difference, initial and boundary conditions! If $P_W(t, n)$ is the analogous quantity, then

$$P_W(0, n) = \begin{cases} 1, & \text{if } n = A \ ; \\ 0, & \text{if } 0 \le n < A \end{cases} \quad ,$$

and the *boundary conditions*

$$P_W(t, A) = \begin{cases} 1, & \text{if } t = A \ ; \\ 0, & \text{if } t < A \end{cases} \quad ,$$

and $P_W(t, 0) \equiv 0$.

Loaded Coin and General Die

For a loaded coin, our recurrence is instead

$$P(t,n) = qP(t-1,n-1) + pP(t-1,n+1) \quad ,$$

and for a loaded die with faces f_i each with probability p_i,

$$P(t,n) = \sum_i p_i P(t-1,n+f_i) \quad .$$

From this we can proceed as before to get linear recurrences for the (binomial) moments, but they are no longer mere polynomials but **exponential-polynomials**, and things get rather messy even for a computer (for an arbitrary die).

Even more generally, we can follow Feller volume 1, section XIV.4, and derive the *probability generating function* in a variable s, say, and to get moments simply differentiate sufficiently many times and plug-in $s = 1$. Note that this does not work for the fair coin case, since then you would get the double root 1 of the indicial equation. It would probably work by using L'Hôspital applied to the general case (of a loaded coin), but I prefer the present approach.

The Maple package RUIN

Everything here (and much more!) is implemented in the Maple package `RUIN` available from the webpage of this article

`http://www.math.rutgers.edu/~zeilberg/mamarim/mamarimhtml/ruin.html` ,

or directly from

`http://www.math.rutgers.edu/~zeilberg/tokhniot/RUIN` .

After you downloaded it into your computer (assuming that you have Maple), stay in the same directory and type

```
read RUIN:
```

and follow the on-line instructions given there. In particular to get the *main* procedures, type `ezra();`, and to get help with any specific procedure, type `ezra(ProcedureName);`. For example, to get help with procedure `StoryW`, type `ezra(StoryW);`.

Some Computer-Generated (Yet Fully Rigorous (Of course!)) and Interesting Theorems

If you type `Story(n,A,x,4);` in `RUIN`, after less than a second, you would get the following output.

Studies in the Classical Gambler's Ruin Problem

by Shalosh B. Ekhad

Consider a player that starts at $x = n$, and with prob. 1/2 walks, at each day, one step to the left or one step to the right, until it either reaches $x = 0$ or $x = A$.

Let D be the random variable: duration of the game. The expectation of D is

$$n(-n + A) \quad .$$

The second moment (about the mean) of D is

$$1/3\, n(-n + A)\left(A^2 - 2\, An + 2\, n^2 - 2\right) \quad ,$$

and setting $n = Ax$, it is

$$-1/3\, A^2 x(x - 1)\left(A^2 - 2\, A^2 x + 2\, A^2 x^2 - 2\right) \quad ,$$

and asymptotically it is, as A goes to infinity

$$-1/3\, x(x - 1)\left(1 - 2\, x + 2\, x^2\right) A^4 \quad .$$

If it starts at the middle, i.e., $x = 1/2$, it is

$$1/24\, A^4 \quad ,$$

which is roughly

$$0.04166666667\, A^4 \quad .$$

The 3^{rd} moment (about the mean) of D is

$$1/15\, n(-n + A)\left(3\, A^4 - 12\, nA^3 + 28\, n^2 A^2 - 10\, A^2 - 32\, An^3 + 20\, An + 16\, n^4 - 20\, n^2 + 4\right) \quad ,$$

and setting $n = Ax$, it is

$$-1/15\, A^2 x(x - 1)\left(3\, A^4 - 12\, A^4 x + 28\, A^4 x^2 - 10\, A^2 - 32\, A^4 x^3 + 20\, A^2 x + 16\, A^4 x^4 - 20\, A^2 x^2 + 4\right) \quad ,$$

and asymptotically it is, as A goes to infinity,

$$-1/15\, x(x - 1)\left(3 - 12\, x + 28\, x^2 - 32\, x^3 + 16\, x^4\right) A^6 \quad .$$

If it starts at the middle, i.e. $x = 1/2$, it is

$$\frac{1}{60}\, A^6 \quad ,$$

which is roughly $0.01666666667\, A^6$

The 4^{th} moment (about the mean) of D is

$$\frac{1}{105}\,n(-n+A)(17\,A^6-67\,nA^5+199\,n^2A^4-84\,A^4-396\,n^3A^3+196\,nA^3+84\,A^2$$

$$-364\,n^2A^2+528\,n^4A^2-28\,An+336\,An^3-396\,An^5+28\,n^2+132\,n^6-168\,n^4+8)\quad ,$$

and setting $n = Ax$ it is

$$-\frac{1}{105}\,A^2x(x-1)(17\,A^6-67\,A^6x+199\,A^6x^2-84\,A^4-396\,A^6x^3+196\,A^4x+84\,A^2-364\,A^4x^2+$$

$$528\,A^6x^4-28\,A^2x+336\,A^4x^3-396\,A^6x^5+28\,A^2x^2+132\,A^6x^6-168\,A^4x^4+8)\quad ,$$

and asymptotically it is, as, A, goes to infinity

$$-\frac{1}{105}\,x\,(x-1)\left(17-67\,x+199\,x^2-396\,x^3+528\,x^4-396\,x^5+132\,x^6\right)A^8\quad .$$

If it starts at the middle, i.e. $x = 1/2$, it is

$$\frac{103}{6720}\,A^8$$

which is roughly $0.01532738095\,A^8$.

The asymptotic Skewness is

$$1/5\,\sqrt{-3\,\frac{\left(3-12\,x+28\,x^2-32\,x^3+16\,x^4\right)^2}{x\,(x-1)\,\left(1-2\,x+2\,x^2\right)^3}}\quad ,$$

and when $x = 1/2$, it is $1/5\,\sqrt{96}$ which is roughly

1.959591794 The asymptotic Kurtosis is

$$-\frac{3}{35}\,\frac{66\,x^4-132\,x^3+99\,x^2-33\,x+17}{\left(1-2\,x+2\,x^2\right)x\,(x-1)}\quad ,$$

and when $x = 1/2$ it is $\frac{309}{35}$ which is roughly 8.828571429.

The sequence of r^{th}-root of r^{th} moment for $x = 1/2$, starting at the second is, divided by A^2 is :

$$[0.2041241452, 0.2554364776, 0.3518576419]\quad .$$

The whole thing took 0.332 seconds of CPU time.

More output

Much more output is available in the webpage of this article, and readers with access to Maple can generate much more on their own.

Apology for the Math Snob

In some sense, the present paper is 'trivial' (modulo known results), since it does not contain genuinely new math (it is mostly in Feller), and is 'just' a description of a computer package. But writing Maple packages implementing even 'trivial' mathematics takes time and is a highly non-trivial endeavor that can't be appreciated by computer-illiterate traditional mathematicians.

Anyway, trivial or not, I was curious about those moments, and couldn't find explicit expressions for them in the literature, so I had to do it myself, and by doing it I illustrated my methodology of *computer generated* mathematics. So in the sense of being a *case-study*, this article is highly non-trivial (if I do say so myself), since it illustrates *by example*, a whole research methodology of experimental yet rigorous mathematics, and with more and more work of this kind, math will become much less trivial than it is now. **Amen**.