

AN EXPERIMENTAL MATHEMATICS PERSPECTIVE ON THE OLD, AND STILL OPEN, QUESTION OF WHEN TO STOP?

LUIS A. MEDINA AND DORON ZEILBERGER

“You got to know when to hold em, know when to fold em, know when to walk away...”

-Kenny Rogers

Supporting Software: This article is accompanied by Maple packages `ChowRobbins`, `STADJE`, and `WALKSab`, and Mathematica packages `Builder.m` (and notebook `Builder.nb`) as well as `STADJE.m`, available from the webpage of this article

<http://www.math.rutgers.edu/~zeilberg/mamarim/mamarimhtml/stop.html>.

1. WHEN TO STOP?

In a delightful and insightful recent “general” article [4], the great probabilist and master expositor Theodore Hill described, amongst numerous other intriguing things, a more than forty-year-old open problem, due to Y.H. Chow and Herbert Robbins [2] that goes as follows:

Toss a fair coin repeatedly and stop whenever you want, receiving as a reward the average number of heads accrued at the time you stop. If your first toss is a head, and you stop, your reward is 1 Krugerrand. Since you can never have more than 100 percent heads, it is clearly optimal to stop in that case. If the first toss is a tail, on the other hand, it is clearly best not to stop, since your reward would be zero...

Then Ted Hill goes on to comment that if the first toss is a tail and the second is a head, then it is good to **go**, since by the law of large numbers, you would eventually do (at least slightly) better than one half. [It turns out that in this case of one head and one tail, the expected gain of continuing the game is larger than 0.6181].

Hill further claims that it is optimal to **stop** if the initial sequence is tail-head-head. [This is *wrong*. It turns out, thanks to our computations, that it is optimal to **go**, and the expected gain is > 0.6693 rather than $2/3$.]

The exact stopping rule, i.e. the decision whether to **stop** or **go**, is still an open problem for (infinitely) many cases. As we will see, it is easy (with computers!) to prove that it is optimal to **go** for many cases where this is indeed the case, but proving *rigorously* that for a given position it is optimal to **stop** is a challenging, still open, problem. It is analogous to disproving vs. proving a mathematical conjecture. To disprove it, all you need is to come-up with a specific counterexample, whereas to prove it, you need to show that no counterexample exists.

2. THE CONTINUOUS LIMIT

Way back in the mid sixties, this problem was tackled by such luminaries as Chow and Robbins themselves [2], Aryeh Dvoretzky [3], and Larry Shepp [6]. Chow and Robbins proved the existence of a stopping sequence, β_n , such that you stop as soon the number of heads minus the number of tails, after n tosses, is $\geq \beta_n$. While Chow and Robbins only proved the *existence* of the “stopping sequence”, Dvoretzky [3] proved that β_n/\sqrt{n} lies between two constants, for n sufficiently large, while Larry Shepp [6] went further and proved that

$$(2.1) \quad \lim_{n \rightarrow \infty} \frac{\beta_n}{\sqrt{n}}$$

exists and equals $0.83992\dots$, a root of a certain transcendental equation.

But this beautiful work, like most of “modern” probability theory, is *asymptotic*, talking about large n . It tells us nothing, for example, about the still open β_8 (presumably 2) and not even about β_{100} . For example, the still-open question whether $\beta_8 = 2$ can be phrased as follows.

*If currently you have five heads and three tails, should you **stop**?*

If you **stop**, you can *definitely* collect $5/8 = 0.625$, whereas if you keep **going**, your expected gain is > 0.6235 , but no one currently knows to *prove* that it would not eventually exceeds $5/8$ (even though this seems very unlikely, judging by numerical heuristics).

3. THE ROLE OF COMPUTERS IN PURE MATHEMATICAL RESEARCH

We really enjoyed Hill’s fascinating article, but we beg to differ on one (important!) issue. Hill ([4], p. 131) claims that:

*“Computers were not useful for solving that problem. In fact, all the problems described in this article were solved using traditional mathematicians’ tools—working example after example with paper and pencil; settling the case for two, three, and then four unknowns; looking for patterns; waiting for the necessary **Aha!** insights; and then searching for formal proofs in each step.”*

So far, this is all *factual*, so there is nothing to disagree with. Ted Hill was merely describing how he and his colleagues do research in pure mathematics. But then came an *opinion* that we do **not** agree with:

“Computers are very helpful for after-the-fact applications of many results, such as backward induction. But in theoretical probability, computers often do not significantly aid the discovery process.”

This may have been true in the *past*, and to a large extent still at *present*, but we believe that in the *future* computers will be more and more useful even—and perhaps especially—in theory, since in addition to their obvious role as **number-crunchers**, they are also starting to do a great job as **symbol-crunchers**, and even as **idea-crunchers**. One recent example is [11], and the present article is another illustration, even though we do quite a bit of number-crunching as well.

4. THE BACKWARD INDUCTION ALGORITHM

The reason that it is so hard to decide (in some cases, for example with 5 heads and 3 tails) whether to **stop** (and collect, *for sure*, the current number of heads divided by the current number of tosses [i.e. $h/(h+t)$]), or to keep **going**, (*expecting* to do better), is the somewhat unrealistic *assumption* that **we live for ever**. Since in *real* life, we eventually would have to quit playing after N tosses, for some finite N , and collect whatever we get then. So let's consider the *bounded* case where the number of coin-tosses is $\leq N$, for a fixed, possibly large, yet finite N . Compromising however with our immortality fantasy, we will let the player collect $1/2$, once reaching the N -th coin toss, if the number of tails exceeds the number of heads, citing the law of large numbers that “guarantees” that “eventually” we will be able to (at least) break even. In other words, we let people who die in debt take advantage of the law of large numbers down in hell. [It turns out that, as far as the soon-to-be-defined limit, $F(h, t)$ goes, one does not need this assumption, and it is possible to insist that the player collects h/N no matter what, but the breaking-even assumption considerably accelerates the convergence.]

Let's call $f_N(h, t)$ the expected pay-off in this bounded game, if you currently have h heads and t tails. Following Chow and Robbins, there is a simple backward induction (dynamical programming) algorithm for computing $f_N(h, t)$ for *all* (h, t) with $h+t \leq N$.

Boundary conditions: when $h+t = N$:

$$(4.1) \quad f_N(h, N-h) = \max(1/2, h/N) \quad , \quad (0 \leq h \leq N).$$

Backward Induction:

$$(4.2) \quad f_N(h, t) = \max \left(\frac{f_N(h+1, t) + f_N(h, t+1)}{2}, \frac{h}{h+t} \right).$$

[If you keep going, the expected gain is $[f_N(h+1, t) + f_N(h, t+1)]/2$, if you stop the expected (and actual) gain is $h/(h+t)$.]

[$f_N(h, t)$ is implemented in procedure `CR(h,t,N)` in `ChowRobbins`. `CRm(h,t,N)` is a faster version].

It is obvious that, for each specific h and t , $f_N(h, t)$ is an increasing sequence in N , bounded above by 1, so we know that the limit

$$(4.3) \quad F(h, t) := \lim_{N \rightarrow \infty} f_N(h, t) \quad ,$$

“exists”.

Fantasizing that we *actually know* the values of $F(h, t)$, (as opposed to knowing that they “exist”), we can decide whether to **stop** or **go**. If $F(h, t) = h/(h+t)$ then we **stop**, and otherwise we **go**. This assumes that the player merely evaluates situations by *expectation*. As we know from the St. Petersburg paradox, expectation is not everything, and a player may choose to *guarantee* collecting $h/(h+t)$ rather than taking a huge chance of eventually getting less. We will later describe other criteria for stopping.

Julian Wiseman [9] estimates $F(0, 0)$ to be 0.79295350640... .

The difficulty in *proving*, for a given number of heads and tails, (h, t) , that it is optimal to stop is that we need *rigorous* non-trivial (i.e. < 1) upper bounds valid

[[36, 31], 105], [[37, 31], 275], [[37, 32], 105], [[38, 32], 235], [[38, 33], 105], [[39, 33], 211], [[39, 34], 105],
 [[40, 34], 193], [[40, 35], 105], [[41, 35], 181], [[41, 36], 105], [[42, 36], 171], [[42, 37], 105], [[43, 37], 165],
 [[43, 38], 107], [[44, 38], 159], [[45, 38], 1039], [[45, 39], 155], [[46, 39], 679], [[46, 40], 153],
 [[47, 40], 513], [[47, 41], 151], [[48, 41], 419], [[48, 42], 149], [[49, 42], 361], [[49, 43], 147], [[50, 43], 321],
 [[50, 44], 147], [[51, 44], 293], [[51, 45], 147], [[52, 45], 271], [[52, 46], 145], [[53, 46], 255],
 [[53, 47], 145]].

7. MORE STATISTICAL INFORMATION

The above strategy for deciding when to stop is entirely based on *expectation*. Even if we pursue this strategy, it would be nice to have more detailed information, like the *standard deviation*, *skewness*, *kurtosis* and even higher moments. Ideally, we would like to know the *full probability distribution*.

Let's call $G_N(h, t; x)$ the fractional polynomial in the variable x (i.e. a linear combination of powers x^a with a rational numbers) such that the coeff. of x^a is the probability of getting exactly a as pay-off in our game, still pursuing the strategy of maximizing the expected gain. Of course $G_N(h, t; 1) = 1$ and $\frac{d}{dx}G_N(h, t; x)|_{x=1} = f_N(h, t)$. We have:

Boundary conditions: when $h + t = N$:

$$(7.1) \quad G_N(h, N - h; x) = x^{\max(1/2, h/N)} \quad (0 \leq h \leq N) \quad .$$

Backward Induction:

$$(7.2) \quad G_N(h, t; x) = \begin{cases} x^{h/(h+t)}, & \text{if } (h, t) \text{ is STOP} \\ \frac{G_N(h+1, t; x) + G_N(h, t+1; x)}{2}, & \text{if } (h, t) \text{ is GO.} \end{cases} ,$$

[$G_N(h, t; x)$ is implemented in procedure CRt(h, t, N, x) in ChowRobbins.]

Once we have $G_N(h, t; x)$, we can easily get all the desired statistical information.

8. ANOTHER WAY TO GAMBLE

In real life we don't always want to maximize our expected gain. Often we have a certain *goal*, let's call it g , and achieving or exceeding it means everlasting happiness, while getting something less would mean eternal misery. In that case we need a different gambling strategy, that is really straightforward. Keep playing until $h/(h+t) \geq g$, and if and when you reach it, **stop**. Otherwise keep going to the end, until $h+t = N$. In that case, of course, the **stop** states are those for which $h/(h+t) \geq g$. It is still of interest to to know what *is* the probability of happiness. Let's call this quantity $P_N(g; h, t)$. We obviously have:

Boundary conditions: when $h + t = N$:

$$(8.1) \quad P_N(g; h, N - h) = \begin{cases} 0, & \text{if } h/N < g \\ 1, & \text{if } h/N \geq g. \end{cases} ,$$

Backward Induction: When $h+t < N$, $P_N(g; h, t)$ equals 1 if $h/(h+t) \geq g$ while it equals $(P_N(g; h+1, t) + P_N(g; h, t+1))/2$ otherwise.

We leave it to the reader to formulate the backward induction scheme for finding the probability generating function for the present strategy.

9. COMPARATIVE GAMBLING

Let's compare the two strategies using *both* criteria. Of course the first one always is better in the *maximum expectation category* and the second is always better in *maximizing the probability of achieving the goal*.

With $N = 200$, at the very beginning, your expected gain, under the first way is 0.7916879464, but your probability

- of getting ≥ 0.6 is 0.6917238235 (the second way gives you probability 0.7753928313, but your expected gain is only 0.6742902054)
- of getting ≥ 0.7 is 0.5625000000 (the second way gives you probability 0.6075176458, but your expected gain is only 0.5787939263)

Much more data can be found by using procedure `SipurCG` in the Maple package `ChowRobbins`, and posted in the webpage of this article.

10. PROBABILITIES OF ESCAPE

The second strategy gives rise to the following interesting computational question:

Fix $a > b \geq 1$ relatively prime. What is the probability that the number of heads divided by the number of tails

- (i) will ever exceed a/b ?
- (ii) will either exceed or be equal to a/b ?

This question was raised and answered by Wolfgang Stadje [8] who proved that this quantity is a root of a certain algebraic equation. A related problem is treated by Nadeau [5].

Stadje's result can also be deduced from the more general treatment by Ayer and Zeilberger [1], that contains a Maple package that automatically derives the algebraic equation for any general set of steps. For practical purposes, however, we found it easiest to compute these probabilities directly, in terms of the discrete functions $W(x, y)$ and $W_s(x, y)$ that count the number of lattice walks from the origin to (x, y) staying in the required region. This is contained in the Maple package `STADJE`.

Here is some data gotten from `STADJE`. The numbers below answer questions (i) and (ii) above, respectively, for each of the listed pairs (a, b) .

- $(a, b) = (2, 1) : 0.6180339887, 0.6909830056 ;$
- $(a, b) = (3, 1) : 0.5436890127, 0.5803566224 ;$
- $(a, b) = (3, 2) : 0.7481518342, 0.7754441182 ;$
- $(a, b) = (4, 1) : 0.5187900637, 0.5362190123 ;$
- $(a, b) = (4, 3) : 0.8091410707, 0.8229424412 ;$
- $(a, b) = (5, 1) : 0.5086603916, 0.5170258817 ;$
- $(a, b) = (5, 2) : 0.5876238826, 0.5996923731 ;$

$$\begin{aligned}(a, b) = (5, 3) &: 0.7158769909, 0.7276461121; \\ (a, b) = (5, 4) &: 0.8453136528, 0.8534748833;\end{aligned}$$

Also of interest is the sequence enumerating the number of walks, staying in the region $y \geq a/bx$, from the origin to a point of the form (n, n) , whose asymptotics can be proved to be of the form $C_1(a, b)4^n/\sqrt{n}$, for some constant $C_1(a, b)$, and the sequence enumerating the number of walks, still staying in the same region, ending at (an, bn) , whose asymptotics has the form $C_2(a, b)((a+b)^{a+b}/(a^a b^b))^n/n^{3/2}$. The Maple package STADJE (and Mathematica package STADJE.m) computes any desired number of terms, and estimates $C_1(a, b)$, $C_2(a, b)$. The webpage of this article contains some sample output.

11. FROM NUMBER-CRUNCHING TO SYMBOL CRUNCHING

So far, we have designed *numerical* computer programs whose outputs were *numbers*. But what about *closed form*? It would be too much to hope for an explicit formula for $f_N(h, t)$ valid for *arbitrary* N, h, t , *but*, with *experimental-yet-rigorous* mathematics, we can find *explicit expressions*, as *rational functions* in n for

$$(11.1) \quad f_{2n+1}(n + \alpha, n - \alpha - m + 1),$$

where n and m are positive integers and α is an integer.

Let

$$(11.2) \quad F(m, \alpha, n) = f_{2n+1}(n + \alpha, n - \alpha - m + 1)$$

for n, m , and α as before. Since $h + t < 2n + 1$, then $F(m, \alpha, n)$ are values below the topmost diagonal on the backward induction triangle.

Some values of $F(m, \alpha, n)$ are not hard to get. For instance, the value of $F(m, \alpha, n)$, for $\alpha \geq 1$ and $1 \leq m \leq 2n$, is given by

$$(11.3) \quad F(m, \alpha, n) = \frac{n + \alpha}{2n - m + 1},$$

whereas the value of $F(m, \alpha, n)$, for $\alpha \leq -m$ and $1 \leq m \leq 2n$, is given by

$$(11.4) \quad F(m, \alpha, n) = \frac{1}{2}.$$

Both formulas can be proved by induction. Hence, we are reduced to finding formulas for $F(m, \alpha, n)$ when $-m < \alpha < 1$.

Our first approach is to make the computer conjecture closed forms for $F(m, \alpha, n)$. For this, we programmed a Mathematica function called **GF** [this function can be found in the webpage of this article]. It takes as input a positive integer m and two variables n and α , and another positive integer *bound*. Here, the computer makes the assumption that $n \geq \text{bound}$. For the guessing part, **GF** uses the auxiliary function **GuessRationalFunction**. This procedure is similar to **GuessRat**, which accompanied the article [7] and can be found in [10]. The output of **GF**, which is the guess formula for $F(m, \alpha, n)$, is a piecewise rational function of n with $m + 2$ pieces.

Example 11.1. For $m = 2$ and $n \geq 3$, GF conjectures

$$(11.5) \quad F(2, \alpha, n) = \begin{cases} 1/2 & \alpha \leq -2 \\ \frac{8n+5}{16n+8} & \alpha = -1 \\ \frac{8n^2+9n+2}{16n^2+8n} & \alpha = 0 \\ \frac{n+\alpha}{2n-1} & \alpha \geq 1 \end{cases}$$

We point out that formulas conjectured by GF only work for n sufficiently large. In fact, empirical evidence suggests that the bound on n grows exponentially in m i.e. as we go down on the backward induction triangle, the bound for which the formulas are valid grows exponentially. As a result, these formulas are not *directly* useful for determining **stop** vs. **go** status.

It is possible to study the recursion formula of $f_n(h, t)$ to get explicit formulas for $F(m, \alpha, 2n+1)$. For example, a simple analysis gives

$$(11.6) \quad F(1, \alpha, 2n+1) = \begin{cases} 1/2 & \alpha \leq -1 \\ \frac{4n+3}{8n+4} & \alpha = 0 \\ \frac{n+\alpha}{2n} & \alpha \geq 1 \end{cases}$$

which is true for $n \geq 1$, and

$$(11.7) \quad F(2, \alpha, 2n+1) = \begin{cases} 1/2 & \alpha \leq -2 \\ \frac{8n+5}{16n+8} & \alpha = -1 \\ \frac{8n^2+9n+2}{16n^2+8n} & \alpha = 0 \\ \frac{n+\alpha}{2n-1} & \alpha \geq 1 \end{cases}$$

which is true for $n \geq 3$. However, these calculations become tedious rapidly.

To our surprise, it turns out that Mathematica, via the built-in functions **Assuming** and **Refine**, is able to handle these recursions and get the desired formulas. We programmed a Mathematica function called **BUILDER**, whose input is an integer m and two variables n and α . **BUILDER** calculates closed-form formulas for $F(m, n, \alpha)$ and provides the smallest n where they start to hold. For instance,

$$(11.8) \quad F(5, \alpha, 2n+1) = \begin{cases} 1/2 & \alpha \leq -5 \\ \frac{64n+33}{128n+64} & \alpha = -4 \\ \frac{32n^2+20n+1}{64n^2+32n} & \alpha = -3 \\ \frac{64n^3+30n^2-13n-3}{64n^4+8n^3-46n^2-5n+3} & \alpha = -2 \\ \frac{128n^3-32n}{256n^5-124n^4-340n^3+91n^2+75n-6} & \alpha = -1 \\ \frac{128n^4-128n^3-32n^2+32n}{512n^5-1280n^4+640n^3+320n^2-192n} & \alpha = 0 \\ \frac{n+\alpha}{2n-4} & \alpha \geq 1 \end{cases}$$

was calculated by **BUILDER** and holds for $n \geq 102$.

The starting places, for n , where the formulas of $F(m, \alpha, n)$ begin to hold, with $1 \leq m \leq 16$, are: 1, 3, 12, 37, 102, 263, 648, 1545, 3594, 8203, 18444, 40973, 90126, 196623, 426000, and 917521 respectively. These values seems to satisfy the recurrence defined by

$$\begin{aligned} a_1 &= 1 \\ a_m &= 2a_{m-1} + r_m \text{ valid for } m \geq 1, \end{aligned}$$

where r_m is given by

$$\begin{aligned} r_1 &= 0 \\ r_2 &= 1 \\ r_3 &= 6 \\ r_m &= 2r_{m-1} + m - 3 \text{ valid for } m \geq 3. \end{aligned}$$

We are pleased to report that the formulas conjectured by GF and the ones found by BUILDER agree.

Acknowledgments. We wish to thank Theodore Hill for very useful Email correspondence. The work of the second author was supported in part by the USA National Science Foundation.

July 23, 2009

REFERENCES

- [1] Arvind Ayer and Doron Zeilberger. Two Dimensional Directed Lattice Walks with Boundaries. *Tapas in Experimental Mathematics, Contemporary Mathematics, Tewodros Amdeberhan and Victor Moll, eds*, 457:1–20, 2008.
- [2] Y.S. Chow and Herbert Robbins. On optimal stopping rule for s_n/n . *Ill. J. Math.*, 9:444–454, 1965.
- [3] Aryeh Dvoretzky. Existence and properties of certain optimal stopping rules. *Proc. Fifth Berkeley Symp. Math. Statist. Prob.*, 1:441–452. Univ. of California Press, 1967.
- [4] Theodore P. Hill. Knowing when to stop. *Amer. Sci.*, 97:126–133, 2009.
- [5] Philippe Nadeau. Walks reaching a line. *DMTCS Proceedings, 2005 European Conference on Combinatorics, Graph Theory and Applications (EuroComb '05)*.
- [6] Larry A. Shepp. Explicit solutions to some problems of optimal stopping. *The Annals of Mathematical Statistics*, 40:993–1010, 1969.
- [7] Andrew Sills and Doron Zeilberger. Disturbing the Dyson conjecture (in a GOOD way). *Experimental Mathematics*, 15:966–990, 1985.
- [8] Wolfgang Stadje. The maximum average gain in a sequence of Bernoulli trials. *Amer. Math. Monthly*, 115, 902–910, 2008.
- [9] Julian D.A. Wiseman. The expected value of $s_n/n \approx 0.79295350640$. <http://www.jdawiseman.com/papers/easymath/coin-stopping.html>
- [10] Doron Zeilberger. The Maple Package GuessRat <http://www.math.rutgers.edu/~zeilberg/tokhniot/GuessRat>
- [11] Doron Zeilberger. The Automatic Central Limit Theorems Generator (and Much More!). *Advances in Combinatorial Mathematics: in honor of Georgy Egorychev, I.Kotsireas and E.Zima, eds. Springer Verlag*, pages 163–172, 2009.

DEPARTMENT OF MATHEMATICS, RUTGERS UNIVERSITY, PISCATAWAY, NJ 08854
E-mail address: lmedina@math.rutgers.edu

DEPARTMENT OF MATHEMATICS, RUTGERS UNIVERSITY, PISCATAWAY, NJ 08854
E-mail address: zeilberg@math.rutgers.edu