Alice and Bob Coin Flip Game

Enrique Treviño

May 4, 2024

1 Puzzle

Daniel Litt introduced the following Puzzle on Twitter on March 16, 2024[1]:

Flip a fair coin 100 times—it gives a sequence of heads (H) and tails (T). For each HH in the sequence of flips, Alice gets a point; for each HT, Bob does, so e.g. for the sequence THHHT Alice gets 2 points and Bob gets 1 point. Who is most likely to win?

2 Exploring a smaller case

Let's explore the case with 4 coin flips. There are 16 possibilities:

1. HHHH. In this case, Alice gets 3 points and Bob gets 0 points. **Alice wins**.
2. HHHT. In this case, Alice gets 2 points and Bob gets 1 point. **Alice wins**.
3. HHTH. In this case, Alice gets 1 point and Bob gets 1 point. **Alice and Bob tie**.
4. HTHH. In this case, Alice gets 1 point and Bob gets 1 point. **Alice and Bob tie**.
5. THHH. In this case, Alice gets 2 points and Bob gets 0 points. **Alice wins**.
6. HHTT. In this case, Alice gets 1 point and Bob gets 1 point. **Alice and Bob tie**.
7. HTHT. In this case, Alice gets 0 points and Bob gets 2 points. **Bob wins**.
8. THHT. In this case, Alice gets 1 point and Bob gets 1 point. **Alice and Bob tie**.
9. HTTH. In this case, Alice gets 0 points and Bob gets 1 point. **Bob wins**.
10. THTH. In this case, Alice gets 0 points and Bob gets 1 point. **Bob wins**.
11. TTHH. In this case, Alice gets 1 point and Bob gets 0 points. **Alice wins**.
12. HTTT. In this case, Alice gets 0 points and Bob gets 1 point. **Bob wins**.
13. THTT. In this case, Alice gets 0 points and Bob gets 1 point. **Bob wins**.
14. TTHT. In this case, Alice gets 0 points and Bob gets 1 point. **Bob wins**.
15. TTTH. In this case, Alice gets 0 points and Bob gets 0 points. **Alice and Bob tie**.
16. TTTT. In this case, Alice gets 0 points and Bob gets 0 points. **Alice and Bob tie**.

Therefore, Bob wins 6 times, Alice wins 4 times and they draw 6 times. We conclude that for four coin flips, the probabilities are

Winner	Probability
Alice	25%
Bob	37.5%
Tie	37.5%

[1] https://twitter.com/littmath/status/1769044719034647001

3 Simulating to get an Answer

Let's now run a simulation where we flip 100 coins, 100000 times and calculate how many times each person won.

Using the Python library *random*, we can write the following code to simulate the 100 coin flips. It will return the difference in points between Alice and Bob.

```
import random

def AliceBobGame():
    currentCoin=random.randint(0,1)
    alice=0
    bob=0
    for i in range(99):
        nextCoin = random.randint(0,1)
        if currentCoin==0 and nextCoin==0:
            alice=alice+1
        elif currentCoin==0 and nextCoin==1:
            bob=bob+1
        currentCoin=nextCoin
    return alice-bob
```

We can write code to simulate the game n times as follows:

```
def sim(n):
    alice=0
    bob=0
    for i in range(n):
        outcome = AliceBobGame()
        if outcome>0:
            alice = alice+1
        elif outcome < 0:
            bob = bob+1
    tie = n - alice-bob
    return alice, bob, tie
```

Running the problem for $n = 100000$, we got the following results:

Winner	Count	Probability
Alice	45703	45.7%
Bob	48513	48.5%
Tie	5784	5.8%

4 Rolling our sleeves and computing exactly

Writing the code above was easy and it gave us strong evidence that Bob will win more often. However, we can actually calculate the probabilities exactly using matrices.

Before we embark on calculating it for 100 coin flips, let's revisit the 4-flip example, but now we will do it using matrices to get the answers instead of going through all 16 cases. The process will be more complicated than simply checking 16 cases, but the advantage is that these ideas will scale and will allow us to calculate the 100 coin flip problem.

4.1 The 4-coin flips case revisited

The idea is to think of the possible states that can happen when a coin flip lands. We could have H in a tie, H with Alice ahead by 1, H with Alice ahead by 2, H with Alice ahead by 3, H with Alice down by 1, etcetera. We can also have T with a tie, Alice ahead by some points or Bob ahead by some points. We will label these states as H_s, T_s where $s \in \{-3, -2, -1, 0, 1, 2, 3\}$ and they represent whether the last coin is H

or T, and s represents the score for Alice (a positive number means she's winning, a negative number means Bob is winning). We can consider that the state before the first flip is T_0 because there won't be any points after one flip, so if you get H you land on H_0 and if you get T you land on T_0. For example, the coin flip sequence $HHTT$ has the following movement between states $T_0 \to H_0 \to H_1 \to T_0 \to T_0$, so it ends at tails and in a tie.

We can now build an adjacency matrix that represents the probability from going from one state to another. Our ordering will be

$$T_0, H_0, T_1, H_1, T_{-1}, H_{-1}, T_2, H_2, T_{-2}, H_{-2}, T_3, H_3, T_{-3}, H_{-3}.$$

From T_0 we can land on H_0 or T_0 each with probability 1/2 and everywhere else, the probability is 0, so the first row of our matrix will be

$$\{\frac{1}{2}, \frac{1}{2}, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0\}$$

.

In general from T_s you go to either H_s or T_s with probability 1/2 and to the other states with probability 0 because if you have TH or TT nobody gets a point so the score stays the same (that is why s didn't change). In the case of H_s, you can get to H_{s+1} with probability 1/2 (when it lands Heads) or you can get to T_{s-1} with probability 1/2 (when it lands Tails). The last four rows are numbers you can't move on from, so we just write a 1 in the same column and 0's elsewhere. The following table suggests shows the probabilities from going from one state to another:

	T_0	H_0	T_1	H_1	T_{-1}	H_{-1}	T_2	H_2	T_{-2}	H_{-2}	T_3	H_3	T_{-3}	H_{-3}
T_0	$\frac{1}{2}$	$\frac{1}{2}$	0	0	0	0	0	0	0	0	0	0	0	0
H_0	0	0	0	$\frac{1}{2}$	$\frac{1}{2}$	0	0	0	0	0	0	0	0	0
T_1	0	0	$\frac{1}{2}$	$\frac{1}{2}$	0	0	0	0	0	0	0	0	0	0
H_1	$\frac{1}{2}$	0	0	0	0	0	0	$\frac{1}{2}$	0	0	0	0	0	0
T_{-1}	0	0	0	0	$\frac{1}{2}$	$\frac{1}{2}$	0	0	0	0	0	0	0	0
H_{-1}	0	$\frac{1}{2}$	0	0	0	0	0	0	$\frac{1}{2}$	0	0	0	0	0
T_2	0	0	0	0	0	0	$\frac{1}{2}$	$\frac{1}{2}$	0	0	0	0	0	0
H_2	0	0	$\frac{1}{2}$	0	0	0	0	0	0	0	0	$\frac{1}{2}$	0	0
T_{-2}	0	0	0	0	0	0	0	0	$\frac{1}{2}$	$\frac{1}{2}$	0	0	0	0
H_{-2}	0	0	0	0	0	$\frac{1}{2}$	0	0	0	0	0	0	$\frac{1}{2}$	0
T_3	0	0	0	0	0	0	0	0	0	0	1	0	0	0
H_3	0	0	0	0	0	0	0	0	0	0	0	1	0	0
T_{-3}	0	0	0	0	0	0	0	0	0	0	0	0	1	0
H_{-3}	0	0	0	0	0	0	0	0	0	0	0	0	0	1

We translate this information into the following matrix:

$$M = \begin{pmatrix}
\frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
\frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{2} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{2} & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1
\end{pmatrix}.$$

This matrix represents the probability to go from one state to another. Note that for the last four rows we just keep the states in the place they already are, as you can't get higher scores with 4 coin flips. If we take the fourth power of M, then we will get a matrix where each entry represents the probability of going from one state to the other in 4 coin flips.

We have

$$M^4 = \begin{pmatrix}
\frac{3}{16} & \frac{3}{16} & \frac{1}{16} & \frac{1}{16} & \frac{3}{16} & \frac{1}{8} & 0 & \frac{1}{16} & \frac{1}{16} & 0 & 0 & \frac{1}{16} & 0 & 0 \\
\frac{1}{16} & \frac{1}{8} & \frac{1}{16} & \frac{3}{16} & \frac{3}{16} & \frac{1}{16} & 0 & 0 & \frac{1}{8} & \frac{1}{16} & 0 & \frac{1}{8} & 0 & 0 \\
\frac{3}{16} & \frac{1}{8} & \frac{3}{16} & \frac{3}{16} & \frac{1}{16} & 0 & 0 & \frac{1}{16} & 0 & 0 & 0 & \frac{3}{16} & 0 & 0 \\
\frac{3}{16} & \frac{1}{16} & \frac{1}{16} & \frac{1}{8} & \frac{1}{8} & \frac{1}{16} & 0 & \frac{1}{8} & 0 & 0 & 0 & \frac{1}{4} & 0 & 0 \\
\frac{1}{16} & \frac{1}{16} & 0 & \frac{1}{16} & \frac{3}{16} & \frac{3}{16} & 0 & \frac{1}{16} & \frac{3}{16} & \frac{1}{8} & 0 & 0 & \frac{1}{16} & 0 \\
\frac{1}{16} & \frac{3}{16} & \frac{1}{16} & 0 & \frac{1}{16} & \frac{1}{8} & 0 & 0 & \frac{3}{16} & \frac{1}{16} & 0 & \frac{1}{16} & \frac{3}{16} & 0 \\
\frac{1}{16} & 0 & \frac{3}{16} & \frac{1}{8} & 0 & 0 & \frac{1}{16} & \frac{1}{8} & 0 & 0 & 0 & \frac{7}{16} & 0 & 0 \\
\frac{1}{8} & \frac{1}{16} & \frac{1}{8} & \frac{1}{16} & 0 & 0 & 0 & \frac{1}{16} & 0 & 0 & 0 & \frac{9}{16} & 0 & 0 \\
0 & \frac{1}{16} & 0 & \frac{1}{16} & \frac{1}{16} & \frac{1}{16} & 0 & 0 & \frac{3}{16} & \frac{1}{8} & 0 & 0 & \frac{7}{16} & 0 \\
\frac{1}{16} & 0 & 0 & 0 & \frac{1}{16} & \frac{1}{8} & 0 & \frac{1}{16} & \frac{1}{16} & \frac{1}{16} & 0 & 0 & \frac{9}{16} & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1
\end{pmatrix}.$$

From this matrix we focus on the first row. The entry in the first column tells us that the probability that the last coin flip is T and the score is tied is $\frac{3}{16}$, while the second column tells us that the probability that the last flip is H and the score is tied is $\frac{3}{16}$, so the game is tied with probability $\frac{3}{8} = 37.5\%$, which matches our previous calculations. The probability that Alice wins can be calculated by adding the entries in columns 3,4,7,8,11, and 12. So the probability that Alice wins is

$$\frac{1}{16} + \frac{1}{16} + 0 + \frac{1}{16} + 0 + \frac{1}{16} = \frac{4}{16} = \frac{1}{4} = 25\%.$$

The probability that Bob wins is calculated by adding the entries in the first row and in the columns 5,6, 9, 10, 13, 14. Therefore, the probability that Bob wins is

$$\frac{3}{16} + \frac{1}{8} + \frac{1}{16} + 0 + 0 + 0 = \frac{6}{16} = \frac{3}{8} = 37.5\%.$$

4.2 The 100 coin flips case

The solution follows an identical strategy as the one we made for the 4 coin flips problem, the difference is that instead of a 14×14 matrix, we will need a 398×398 matrix[2] with the states being

$$T_0, H_0, T_1, H_1, T_{-1}, H_{-1}, \ldots, T_{99}, H_{99}, T_{-99}, H_{-99}.$$

Let's call this 398×398 matrix A. Let $A(i,j)$ be the entry in the i-th row and j-th column of A. For $k = 1, \ldots, 99$, the $4k-1$-th state represents T_k, the $4k$-th state represents H_k, the $4k+1$-th state represents T_{-k} and the $4k+2$-th state represents H_{-k}.

The first two rows are:

$$A(1,j) = \begin{cases} \frac{1}{2} & \text{if } j = 1 \text{ or } j = 2 \\ 0 & \text{otherwise} \end{cases},$$

and

$$A(2,j) = \begin{cases} \frac{1}{2} & \text{if } j = 4 \text{ or } j = 5 \\ 0 & \text{otherwise} \end{cases}.$$

Just as when we built M, the last four rows of A are formed by many zeroes and the 4×4 identity for the last 4 columns, so

$$A(i,j) = \begin{cases} 1 & \text{if } j = i \\ 0 & \text{otherwise} \end{cases} \qquad \text{for } 395 \leq i \leq 398.$$

[2]Technically, we can cut a lot of rows and columns by eliminating states that are impossible to reach, for example anything with an index of -51 or lower can be eliminated because Bob can't have more than 50 points in 100 coin flips. The H_{-50} can also be removed because Bob can't have 50 points if the last coin flip is H. Making these changes would reduce the size of the matrix to a 299×299 matrix. However, setting up the matrix with these rows removed would be more complicated because the periodic pattern would be broken.

We know that $T_k \to \{T_k, H_k\}$, so for $i = 4k-1$ or $i = 4k+1$

$$A(i,j) = \begin{cases} \frac{1}{2} & \text{if } j = i \text{ or } j = i+1 \\ 0 & \text{otherwise} \end{cases} \qquad \text{for } i = 4k-1 \text{ or } i = 4k+1 \text{ and } 1 \le k \le 98.$$

The only rows left to account for are when $i = 4k$ and when $i = 4k+2$ for $1 \le k \le 98$. For $i = 4k$, the state we are looking at is H_k. With Heads we get to H_{k+1} and with tails we get to T_{k-1}. So the case $k = 1$ is special (because we get T_0). We have to do the fourth row separately from the other multiples of 4, we get

$$A(4,j) = \begin{cases} \frac{1}{2} & \text{if } j = 1 \text{ or } j = 8 \\ 0 & \text{otherwise} \end{cases}.$$

Now for the other multiples of 4, we have

$$A(4k,j) = \begin{cases} \frac{1}{2} & \text{if } j = 4k-5 \text{ or } j = 4k+4 \\ 0 & \text{otherwise} \end{cases} \qquad \text{for } 2 \le k \le 98.$$

For $i = 4k+2$ we are looking at the state H_{-k}. We can reach H_{-k+1} or T_{-k-1}. Note that our work suggests $H_{-k} \to H_{-k+1}$ would be a change from $4k+2$ to $4k-2$. This even works in the case $k = 1$ (because H_0 is also in the $4 \cdot 1 - 2$ position). Also $H_{-k} \to T_{-k-1}$ is a change from $4k+2$ to $4(k+1)+1 = 4k+5$. Therefore we get

$$A(4k+2,j) = \begin{cases} \frac{1}{2} & \text{if } j = 4k-2 \text{ or } j = 4k+5 \\ 0 & \text{otherwise} \end{cases} \qquad \text{for } 1 \le k \le 98.$$

We have now covered all 398 rows. We can then evaluate A^{398}. The probability that Alice and Bob tie can be evaluated by adding the first two entries in the first row, the probability that Alice wins can be evaluated by adding the entries in the columns $4k-1, 4k$ for $k = 1, 2, \ldots, 99$ of the first row, and the probability that Bob wins can be evaluated by adding the other entries in the first row.

We get the following results:

Winner	Exact Probability	Percentage
Alice	$\frac{1450319873098552085952721068 51}{3169126500570573503741758013 44}$	45.764%
Bob	$\frac{1539665596047401055899223888 55}{3169126500570573503741758013 44}$	48.583%
Tie	$\frac{89570515712310180944906528 19}{1584563250285286751870879006 72}$	5.653%

5 Intuitive Explanation

The short answer on why Bob wins more often than Alice is that on average they both average the same score (by linearity of expectation, the expected score is 99/4), but Alice has more variance than Bob, so Bob is often around the expected score while Alice can overshoot it or undershoot it by a lot. One can think of it as Bob wins many close matches and Alice wins many blowouts (and some close matches).

A very nice explanation that is more rigorous is suggested by Sridhar Ramesh in a Twitter thread written on March 19, 2024[3]. The idea is to imagine a random walk where we take a step up when we see HH and we take a step down when we see HT^nH, however we move at a certain pace. So HH is a move 1 unit up at 1 unit per flip, while and HT^nH is a unit down at a rate of $1/n+1$ per flip, so if the flips end HTT then we didn't move a full unit down, but we moved 2/3 of a unit down. Let's illustrate it with the 4-coin example:

1. HHHH. We move 3 units up and end at 3.
2. HHHT. We move 2 units up followed by 1/2 down to end at 2.5.
3. HHTH. We move 1 unit up and 1 unit own to end at 0.

[3]https://twitter.com/RadishHarmers/status/1770217473716932647

4. HTHH. We move 1 unit down and 1 unit up to end at 0.
5. THHH. We move 2 units up to end at 2.
6. HHTT. We move 1 unit up followed by 2/3 down to end at 1/3.
7. HTHT. We move 1 unit down dollowed by 1/2 down to end at -1.5
8. THHT. We move 1 unit up followed by 1/2 down to end at 0.5
9. HTTH. We move 1 unit down to end at -1.
10. THTH. We move 1 unit down to end at -1.
11. TTHH. We move 1 unit up to end at 1.
12. HTTT. We move 3/4 of a unit down to end at -0.75.
13. THTT. We move 2/3 of a unit down to end at -2/3.
14. TTHT. We move 1/2 of a unit down to end at -0.5.
15. TTTH. We don't move and end at 0.
16. TTTT. We don't move and end at 0.

Note that 6 times we end a walk at a negative number (all Bob's wins!) and 6 times we end at a positive number (4 Alice wins, 2 ties). We also end up at 0 four times (all ties).

If we now consider 100 coin flips, when the walk ends at a number ≥ 1, Alice wins, when the number ends at 0 is a tie, when the number is between 0 and 1 is a tie and when the number is negative Bob wins.

Why does this show us Bob wins? Well, as Ramesh explains, while walking on different speeds, when one returns to the origin, one can invert the steps and also return to the origin, so we can focus on the last move away from the origin and it is equally likely to be above or below. Therefore, the probability of ending above the origin is the same as the probability of ending below the origin in a random walk. However, all the times we end below the origin are wins for Bob, while some of the times we end up above the origin are not Alice's wins. This is why Bob wins more often.

As a remark, note that for the walk to end between 0 and 1, the last coin flip has to be T. Whenever the walk ends at a non-integer, it had as its last coin flip T. Whenever a walk ends at an integer, the last coin flip is H unless it's the $TTT \cdots T$ example (because if there is an H before $T \cdots T$, then the last step moves downwards). Therefore, besides the $T \cdots T$ case, all ties that end in T have a value between 0 and 1 (not inclusive) and all ties that end in H have the random walk end at 0. If it's equally likely for ties to end in H and T, then that means that the difference between Bob's probability of winning and Alice's probability of winning is approximately half the probability of tying. Empirical data shows that it is close to half, but not exactly half, so perhaps there's a slight bias for ties to end with one coin flip over the other.