Hitting 10

You toss a coin repeatedly. If heads is worth 1 and tails is worth 2, and you add up the score as you go along, what is the probability that you'll hit 10?

That is to say, 10 might be one of the sums you get, but it also might not—if you hit 9 and then roll tails, you'll go straight to 11 and miss 10. The question asks, if you do this a large number of times, what proportion of the times will you get 10 as one of your sums?

A note on notation: it makes sense to jettison the terminology "heads" and "tails" and use "1" and "2", as if the coin had a 1 on one side and a 2 on the other. This is what I will do.

As usual, I start out collecting data. Each of 35 students runs 10 trials and keeps track of the number of "successes" (hitting 10). The scores range from 4/10 to 9/10 successes, with an overall success rate of $216/350 = 0.62$. That's an empirical estimate of our probability.

One student asserts right away that the answer should be 2/3. Here is her argument. Think of the possible sequences of sums that go "past" 10. We could have

$$8, 10$$
$$9, 10$$
$$9, 11$$

These are the only possibilities—that is, every sequence will contain one (and only one) of these pairs. Now two of these hit 10 and one misses 10, for a probability of 2/3. An interesting argument.

But a number of students are uncomfortable with it and there is a clear need for some care in the analysis. But where to begin? Someone suggests that if we knew the probability of hitting 8 or 9, then we might be able to figure out the probability of hitting 10. That sounds like induction—working up from small numbers. So let's take some small numbers N and see what we get.

Some notation: **we'll denote by P_N the probability of hitting the target N.**

Why not start at the very beginning. What's the probability of hitting $N=1$? Well, to hit 1, you'll have to toss 1 right away, and that will happen with probability ½. So $P_1 = ½$.

Okay, let's go for $N=2$. Here I get a verbal argument. Think about the first toss: half the time you toss 2 and succeed right away. What about the other half when you toss 1? Well, then you toss again and half the time you toss another 1 (success) and the other half you toss 2 (failure). So with probability ½ + ¼ you hit 2 and with probability ¼ you miss it. So $P_2 = ¾$.

Now on to *N*=3. This time a student comes to the board and tabulates all the possible sequences that would arise, stopping each time when it is known whether 3 has been hit. There are 5 possible sequences, of which 3 are successful, for a probability of 3/5. I let the table stand, and take a vote. Nearly everyone accepts the result. But there are a few abstainers with cross looks and wrinkled brows.

Buoyed up by this show of support, the student carries on. He has done the same for *N*=4 and found 8 possible sequences of which 5 are successes. Now 3, 5 and 8 are successive Fibonacci numbers, and he asserts that it continues in this way. Thus P_N is always the ratio of two successive Fibonacci numbers.

Well, what can I say? Wow. Certainly the class is cowed. "Comments?" I ask crisply. Silence, even from the wrinkled brows. We've encountered the Fibonacci numbers. What could be better? I just has to be right.

To trouble-shoot, I suggest we go back and do the same analysis for the case *N*=2 that we have already done. It shows three possibilities, two of which are successes. Thus P_2 should be 2/3. But we've just shown that $P_2 = 3/4$. So there's certainly a flaw in the reasoning here.

Let's compare this with the original argument—there we noted that the initial 2 (the third row) occurs ½ the time, so this contributes ½ to the probability right away. And the other two cases each contribute ¼. That is, it would seem that the rows of the table should be *weighted* with their probability of occurrence. If we include these weights (as a third column), and add up all the S-weights, we do get the correct answer of ¾.

It would appear that the flaw in the "Fibonacci" calculation is the various rows in the table haven't been weighted by their probability of occurrence. Let's see what happens if we do that with the *N*=3 table. The new weighted table now gives a success probability of 1/8 + 1/4 + 1/4 = 5/8 and this is in fact the correct answer. It also checks out with a couple of approaches that other students have been trying.

Note that since the third column represents the probabilities of the different possible outcomes, they have to sum to 1. That's a nice error check. [We've had to abandon that amusing Fibonacci excursion, but see problem 3.]

The N=3 table	
111	S
112	F
12	S
21	S
22	F

The N=4 table	
1111	S
1112	F
112	S
121	S
122	F
211	S
212	F
22	S

The N=2 table	
11	S
12	F
2	S

The N=2 table with weights		
11	S	¼
12	F	¼
2	S	½

The N=3 table with weights		
111	S	1/8
112	F	1/8
12	S	1/4
21	S	1/4
22	F	1/4

I give the class the job of constructing the weighted table for $N=4$, and we get the answer of 11/16. It starts to get hard to organize the entries in a systematic way. Some students are using tree diagrams instead of tables, and they provide an easy system, but they get large quickly.

The N=4 table with weights		
1111	S	1/16
1112	F	1/16
112	S	1/8
121	S	1/8
122	F	1/8
211	S	1/8
212	F	1/8
22	S	1/4

Can you track the correspondence between the tree diagram and the table?
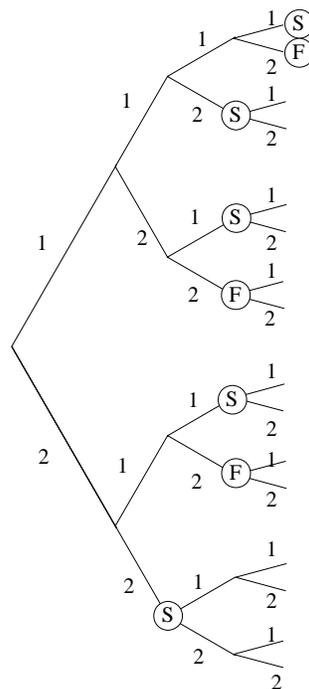
Perhaps it's time to start tabulating the P_N values to see if we can see some useful patterns in the numbers. So far, we have calculated the first 4 values. Any guesses as to what comes next? One clear pattern is seen in the denominators—2, 4, 8, 16—at the Nth stage we get 2^N.

Well, an interesting pattern is observed in the numerators as well: 1, 3, 5 and 11. Add consecutive pairs:

N	P_N
1	1/2
2	3/4
3	5/8
4	11/16
5	?

$1 + 3 = 4$ which is the denominator of 3
$3 + 5 = 8$ which is the denominator of 5
$5 + 11 = 16$ which is the denominator of 11.

Can this be a coincidence? The prediction it gives for P_5 is 21/32. And someone has been constructing the table for $N=5$, and reports that this is correct! Wow. Nice.

Perhaps we should stop now and try to see why this interesting numerical pattern might hold (and thereby get a proof of its validity). But the class is impatient to complete the table—and so that's what we do.

Okay! We have a candidate: $P_{10} = 683/1024$. How can we be sure that it's correct?

N	P_N
1	½
2	¾
3	5/8
4	11/16
5	21/32
6	43/64
7	85/128
8	171/256
9	341/512
10	683/1024

A hand goes up in the back. It is Devon.
It's correct.
I beg your pardon?
It's correct. The answer for P_{10} is correct.
How do you know?
I worked it out.

You worked it out? You calculated P_{10}?
Yes.
How?
Uhmm…
You made a table of all the possibilities?
Not exactly...
A tree?
Not really...

I invite Devon to the front, and give him the chalk and a clean panel. He proceeds to display an extremely elegant accounting of all possible paths leading to a success for $N=10$. With a bit of reorganization on my part, this is the table he puts up.

Different routes to get to 10			
sequence composition	*Prob. of sequence*	*number of sequences*	*total probability*
10 ones	$\frac{1}{2}^{10}$	1	$1/2^{10}$
8 ones & 1 two	$\frac{1}{2}^{9}$	9	$9/2^{9}$
6 ones & 2 twos	$\frac{1}{2}^{8}$	$8!/6!2! = 28$	$28/2^{8}$
4 ones & 3 twos	$\frac{1}{2}^{7}$	$7!/4!3! = 35$	$35/2^{7}$
2 ones & 4 twos	$\frac{1}{2}^{6}$	$6!/2!4! = 15$	$15/2^{6}$
5 twos	$\frac{1}{2}^{5}$	1	$1/2^{5}$

Having put the table on the board, he then calculates the sum:

$$\frac{1}{2^{10}} + \frac{9}{2^{9}} + \frac{28}{2^{8}} + \frac{35}{2^{7}} + \frac{15}{2^{6}} + \frac{1}{2^{5}}$$

$$= \frac{1}{2^{10}}\left[1 + 2\cdot 9 + 4\cdot 28 + 8\cdot 35 + 16\cdot 15 + 32\right] = \frac{683}{2^{10}}$$

A convincing result!

Needless to say, this "full calculation" is not the route I intended the class to follow. Indeed, many of them have not yet encountered the combinatorial coefficients. And I can feel a definite level of anxiety rising from the class.

"What we have here," I say, "Is a complete (and remarkably compact) listing of cases." Such an approach is of great theoretical importance, and you would encounter it in any introductory finite math course. So I'm very happy for you to have seen an example here of what it looks like. However the method has some disadvantages. It is computationally heavy. How well would it work if I had asked for the probability of hitting 100?" I look at Devon and he wrinkles his brow. "And even more—could this approach give us a general formula for P_N?"

In fact in problems of this type there is often an approach which does not require listing all the successful cases, which uses a nice piece of reasoning to get around that. Let's get back to some of those intriguing patterns.

What happened here really blew me away. Devon hadn't wasted any time either—he had in fact obtained the answer before we had completed the table, and was just waiting for the chance to put it forward. It turns out he was taking OAC finite math that very term (the year was 1995), and this problem hit him right at a golden moment.

Most of the class has trouble seeing what Devon has done. It helps to look back at the $N=4$ table. Notice, for example, that there are three rows which have 2 ones and 1 two, each with probability 1/8. Well, in his accounting, these would all be listed on one row and there would be a 3 in his third column to say that he's counted three rows of the "expanded" table. For example, the 28 in column 3 says that there are 28 ways to arrange a string of 6 ones and 2 twos, and so these would have occupied 28 rows in the extended table. So the table at the left gives us a very economical presentation.

By now, the class has noticed more patterns in the P_N table. For example, if you double any numerator, you get the next numerator ±1, where you alternate the + and the − .

N	P_N
1	1/2
2	3/4
3	5/8
4	11/16
5	21/32

The point is that these *numerical* patterns should correspond to *algebraic* relationships between the P_N. And if we can formulate these, we might be able to "see" why they must hold.

Let's take, for example, the original numerator pattern:
$$1 + 3 = 4$$
$$3 + 5 = 8$$
$$5 + 11 = 16, \text{etc.}$$
What does this tell us about relationships among the P_N ? To have a specific case, let's take:
$$5 + 11 = 16.$$

To get the P_N into the picture, divide by 16:

$$\frac{5}{16} + \frac{11}{16} = 1$$

Now the 11/16 is clearly P_4 . But what is the 5/16. Well, it would seem to be half of P_3. So we write this as:

$$\frac{P_3}{2} + P_4 = 1$$

Okay. That's an intriguing relationship. Who can see a reason why it should be true?

Well, since the two terms on the left sum to 1, and the second is the probability of hitting 4, the first *has* to be the probability of *not* hitting 4. So that becomes the question—can we see a reason why P₃/2 should be the probability of not hitting 4?

Well, it takes a minute, but it's all too easy. The crucial observation is that *there's only one way to miss 4 and that's to hit 3 and then toss a 2*. And the chances of doing that are $P_3 \times$ ½, the first to hit the 3, and the second to toss the 2. That's real nice.

This is a general result: *The only way to miss n+1 is to hit n and then toss a 2*. This can be written:

$$1 - P_{N+1} = \frac{1}{2} P_N$$

If we rewrite it as:

$$P_{N+1} = 1 - \frac{1}{2} P_N$$

It is in recursive form and this allows us to recursively calculate all the P_N starting at $P_1 = ½$.

I give the students some time to find this lovely argument on their own, but not many succeed. They seem to find it a bit foreign.

One might have thought that this recursive approach would have been tried at the very beginning of the problem. But for all it's simplicity and directness, it does not seem to occur naturally. It's an approach that must be learned and experienced.

One more question—can we work out P_{100} from this? Of course we could start with P_1 and iterate the recursion 99 times. That would be easy enough with a spreadsheet, but where would it get us? What we'd really like is a *formula* for P_N in terms of N. And that's really a question of *solving* the recursion.

We have what's called a first-order linear inhomogeneous recursion and there are lots of different "methods" around to solve it. But my favorite is the one that I find high school students doing if they're not told what to do, and that's just to iterate the recursion again and again but keeping the terms in expanded form until one sees a pattern.

$$P_{N+1} = 1 - \frac{1}{2} P_N$$

We have a recursive formula for P_N. Can we solve it?

To make the calculations easier (and more transparent!) we write the recursion as

$$P_{N+1} \;=\; 1 + aP_N \qquad\qquad \text{with } a = -\tfrac{1}{2}.$$

Then start at the beginning:

$P_0 = 1$
$P_1 = 1 + aP_0 = 1 + a.$
$P_2 = 1 + aP_1 = 1 + a(1 + a) \;=\; 1 + a + a^2$
$P_3 = 1 + aP_2 = 1 + a(1 + a + a^2) \;=\; 1 + a + a^2 + a^3$
$P_4 = 1 + aP_3 = 1 + a(1 + a + a^2 + a^3) \;=\; 1 + a + a^2 + a^3 + a^4$

The pattern is clear and it is clear that it will continue (a formal proof of that can be made using mathematical induction).

$$P_N = \; 1 + a + a^2 + a^3 + \ldots + a^N \;=\; \frac{1-a^{N+1}}{1-a}$$

Notice what a difference that "a" makes. Using a general parameter there instead of the –½ really allows us to focus on the developing pattern as we move from one N to the next.

using the formula for the sum of a geometric series.
Now putting $a = -\tfrac{1}{2}$:

$$P_N = \; \frac{1-(-1/2)^{N+1}}{1-(-1/2)} \;=\; \frac{1-(-1/2)^{N+1}}{3/2} \;=\; \frac{2}{3}\Big[1-(-1/2)^{N+1}\Big].$$

We have a formula! The probability of hitting 10 is then

$$P_{10} = \; \frac{2}{3}\Big[1-(-1/2)^{11}\Big] \;=\; \frac{2}{3}\Big[1-\frac{-1}{2048}\Big] \;=\; \frac{2}{3}\Big[\frac{2049}{2048}\Big] \;=\; \frac{683}{1024}.$$

as we obtained with our combinatorial calculations.

The general formula is nice. It tells us that the probabilities will be successively above and below 2/3, and will approach 2/3 in the limit. You might have managed to already argue that the probability should get close to 2/3 for large N.

Question. There's something a bit unexpected about the P_N formula and that's the 3 in the denominator. Isn't the denominator supposed to be 2^N? Argue that P_N will always be able to be written as a fraction with 2^N in the denominator.

Why should the P_N approach 2/3?
Here's a nice way of thinking about it. Suppose we make a large number of tosses, say 200. Then we expect 100 1's and 100 2's, so our total sum should be about (on average?) 300. But how many numbers between 1 and 300 will we have hit?—well we get a new hit with every toss, so we'll hit 200 numbers. That means we hit a proportion 200/300 of the numbers less than 300. So the probability of hitting any particular number should be around 2/3.

Problems

1. Harry and Sally play the following game with a biased coin. Sally flips first. If she gets heads she wins; if she gets tails she gives the coin to Harry. Then Harry flips and if he gets heads he wins; if he gets tails he gives the coin to Sally. The game continues until someone flips heads and wins.

(a) If heads comes up with probability 1/4, what is the probability Sally will win?
(b) If over many repetitions of the game, Sally wins 75% of the time, what would you say was the bias on the coin?
(c) Let the bias on the coin be such that heads occurs a proportion x of the time, and suppose Sally wins with probability p. Find simple formulas for p in terms of x and for x in terms of p.

2. Annie and Bill each have a fair die, which they throw alternately. Annie scores if she throws a multiple of three, and Bill scores if he throws a multiple of two. The winner is the first person to score. Show that the game is fair if Annie rolls first.
[Note: If I'm using this problem in class, it works well to put the class into groups of two to play this game a few times, and collect some data. We see if the game really seems to be fair, and we get a "hands-on" feel for the analysis.]

3. Here's a generalization of 2. Suppose that Annie and Bill each spin a pointer with n equally likely outcomes numbered 1 to n. Annie scores if she throws any of the first numbers 1, 2, … k and Bill scores if he throws any of the numbers $k+1$, $k+2$, … n. If Annie goes first, for what numbers n and k is the game fair?

4. A flips a fair coin 100 times, and B flips a fair coin 101 times. What is the probability that B gets more heads than A? [This is an awesome problem and works well with an exploratory approach. Start by solving simpler problems.]

5. *Following Fibonacci.* What the Fibonacci guy discovered (above) is that if you count the *number* of different sequences that you might generate in deciding whether a target n was "hit" you got a Fibonacci number, and if you counted the number of those that were "successful," in that they hit the target, that was the preceding Fibonacci number. In fact, that's a very nice result, not about probabilities, but about numbers of sequences of a certain type. Track these ideas down and write a nice little essay about the subject.

6. Follow the methods developed above to solve the recursion
$$x_{n+1} = 3 - \frac{3}{4} x_n$$
starting with $x_0 = 1$. That is, find a general formula for x_n .