Contents

1

PROBABILITY IN YOUR HEAD

Peter Winkler

Probability theory is a well-developed science, and most probability puzzles, if translated into precise logical form, could in principle be solved by machine. But the theory also provides some remarkable problem-solving tools with which we humans, using our imagination, can often find delightful shortcuts.

Here are eight puzzles that you can try to solve by taking pencil in hand and “doing the math," but each can also be solved in your head, just by reasoning.

1 Problems

1.1 Flying Saucers

A fleet of saucers from planet Xylofon has been sent to bring back the inhabitants of a certain apartment building, for exhibition in the planet zoo. The earthlings therein constitute 11 men and 14 women.

Saucers arrive one at a time and randomly beam people up. However, owing to the Xylofonians’ strict sex separation policy, a saucer cannot take off with humans of both sexes. Consequently, a saucer will continue beaming people up until it acquires a member of a second sex; that human is immediately beamed back down, and the saucer takes off with whomever it already has on board. Another saucer then swoops in, again beaming up people at random until it gets one of a new gender, and so forth, until the building is empty. What is the probability that the last person beamed up is a woman?

1.2 Points on a Circle

Three points are chosen at random on a circle. What is the probability that there is a semicircle of that circle containing all three?

1.3 Meet the Williams Sisters

Some tennis fans get excited when Venus and Serena Williams meet in a tournament. The likelihood of that happening normally depends on seeding and talent, so let us instead assume an idealized elimination tournament of 64 players, each as likely to win as to lose any given match, with bracketing chosen uniformly at random. What is the probability that the Williams sisters get to play each other?

1.4 Service Options

You are challenged to a short tennis match, with the winner to be the first player to win four games. You get to serve first. But there are options for determining the sequence in which the two of you serve:

1. Standard: Serve alternates (you, her, you, her, you, her, you).
2. Volleyball style: The winner of the previous game serves the next one.
3. Reverse volleyball style: The winner of the previous game receives in the next one.

Which option should you choose? You may assume it is to your advantage to serve. You may also assume that the outcome of any game is independent of when the game is played and of the outcome of any previous game.

1.5 Who Won the Series?

Two evenly-matched teams meet to play a best-of-seven World Series of baseball games. Each team has the same small advantage when playing at home. As usual, one team (say, Team A) plays games 1 and 2 at home, and, if necessary, plays games 6 and 7 at home. Team B plays games 3, 4 and, if needed, 5 at home.

You go to a conference in Europe and return to find that the series is over, and six games were played. Which team is more likely to have won the series?

1.6 Random Rice

You go to the grocery store needing 1 cup of rice. When you push the button on the machine, it dispenses a uniformly random amount of rice between nothing and 1 cup. On average, how many times do you have to push the button to get (at least) a cupful?

1.7 Six with No Odds

On average, how many times do you need to roll a die to get a 6, given that you do not roll any odd numbers *en route*? (Hint: The answer is not 3.)

1.8 Getting the Benz

Your rich aunt has died and left her beloved 1955 Mercedes-Benz 300 SL Gullwing to either you or one of your four siblings, according to the following stipulations. Each of the five of you will privately write "1," "2," or "3" on a slip of paper. The slips are put into a bowl to be examined by the estate lawyers, who will award the car to the heir whose number was not entered by anyone else. (If there is no such heir, or more than one, the procedure is repeated.)

For example, if the bowl contains one 1, two 2s and two 3s, the one who put in the 1 gets the car.

After stewing and then shrugging your shoulders, you write a 1 on your slip and put it in the bowl. What the heck, a 1/5 chance at this magnificent vehicle is not to be sneezed at! But just before the bowl is passed to the lawyers, you get a sneak peek and can just make out, among the five slips of paper, one 1 (which may or may not be yours), one 2, and one 3.

Should you be happy, unhappy, or indifferent to this information?

2 Solutions

2.1 Flying Saucers

This puzzle—and many others—is more easily tackled if we rephrase the question, perhaps by putting the randomness up front. Suppose we imagine that the building occupants are first lined up uniformly at random, then picked up by the flying saucers starting from the left end. Then the gender of the rightmost person in line would determine whether the last person picked up is a woman, and of course that gender is female with probability $14/(14+11)=56\%$.

But there's a problem: This model is not correct, because the next saucer's beam-ups are newly randomized and might not begin with the person who was last rejected. You can see the difference by examining the case where there are just two women and one man.

So, let us imagine a different model, where the remaining occupants are re-lined up randomly each time a saucer arrives. Then the *next-to-last* saucer must be facing a line consisting entirely of one or more men followed by one or more women; or, one or more women followed by one or more men. Because these possibilities are equally likely (reversing the line, for example, transforms one set of possibilities to the other), the probability that the last saucer will be greeted by women is 1/2.

2.2 Points on a Circle

This puzzle, a classic, was suggested to me by combinatorics legend Richard Stanley of the Massachusetts Institute of Technology and the University of Miami.

Let us pick the three points in a funny way: Choose three random diameters, that is, three lines through the center of the circle at uniformly random angles. Each diameter has two endpoints on the circle, giving us six points, which we can label A, B, C, D, E, F clockwise around the circle, beginning anywhere. Then

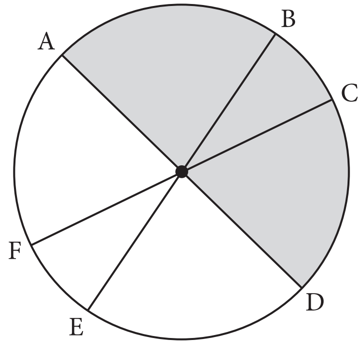

Figure 1.1. Three diameters with their endpoints labeled sequentially. Any three consecutive points, such as A, B, C or B, C, D, are easily seen to lie on a common semicircle. Nonconsecutive points that include one endpoint from each of the three diameters, such as A, C, E, do not lie on a common semicircle.

we will use a coin flip (three times) to decide, for each line, which of these two intersections becomes one of our three points.

It not hard to see that if the chosen points are consecutive (i.e., A, B, C, or B, C, D, etc.), up to F, A, B (six possibilities), then they are contained in a semicircle; otherwise not. This is illustrated in Figure 1.1.

Since there are eight possible outcomes for the coin flips, the probability that the desired semicircle exists is a whopping 3/4. (If you ask your friend to pick three random points on a circle, I expect they are less likely than not to be contained in a semicircle.)

Similar reasoning shows that if n points are chosen, $n \geq 2$, the probability that they are contained in some semicircle is $2n/2^n$. In fact, you can even apply a version of this argument in higher dimensions; for example, to deduce that the probability that four random points on a sphere are contained in some hemisphere is $1/8$.

2.3 Meet the Williams Sisters

This puzzle appears in Frederick Mosteller's wonderful little book, *50 Challenging Problems in Probability* [2], but the solution he offers involves working out examples, then guessing a general solution and proving its correctness by induction. Here, instead, is a solution you can work out in your head.

Given the problem's symmetry conditions, each of the $\binom{64}{2} = 63 \times 32$ pairs of players has the same probability of meeting. Since 63 matches are played (remember that all but one of the 64 players needs to be eliminated to arrive at the winner), the probability that the Williams sisters meet is $63/(63 \times 32) = 1/32$.

Here is another approach that almost works. It takes two coin flips on average to get a head and thus two matches in our random tournament to get a loss. Therefore, a player will meet on average two other players, so Serena's probability of meeting Venus is 2/63. Wait, that is not quite right. Can you find the flaw?

2.4 Service Options

This puzzle was inspired by one I heard from Dick Hess, author of *Golf on the Moon* [1] and other delightful puzzle books.

Assume that you play lots of games (maybe more than is needed to determine the match winner), and let A be the event that of the first four served by you and the first three served by your opponent, at least four are won by you. Then it is easily checked that no matter which service option you choose, you will win if A occurs and lose otherwise. Thus, your choice makes no difference. Notice that the independence assumptions mean that the probability of the event A, since it always involves four particular service games and three particular returning games, does not depend on when the games are played or in what order.

If the game outcomes were not independent, the service option could make a difference. For example, if your opponent is easily discouraged when losing, you might benefit by using the volleyball scheme, in which you keep serving if you win.

The idea that playing extra games may be useful for analytic purposes (despite—no, because!—they do not affect the outcome) will be even more critical in the next solution.

2.5 Who Won the Series?

This nice question came to me from Pradeep Mutalik, who writes the math column for the excellent online science magazine *QUANTA*. The solution is my own.

The key here is that *potential* games are as important as actual ones when it comes to computing the odds. It is tempting to think, for example, that Team A's extra home game is less of a factor, because the series does not usually go to seven games. But this is false reasoning: You *may as well* assume all seven games are played (since it makes no difference to the outcome if they are). Thus, the 4-to-3 advantage in home games enjoyed by Team A is real and is unaffected by the order of the games.

(To see a more extreme example of this phenomenon, imagine the series winner is to be the first team to win 50 games, and that the first 49 are home games for Team B and the rest for Team A. Then, since the series is a big favorite to end before game 98, most games will *probably* be played on Team B's home field—yet, it is Team A that has the advantage. This, again, can be seen by imagining that 99 games are played regardless; perhaps tickets for all 99 have already been sold and the fans don't mind watching games played after the series outcome is decided.)

Similar reasoning shows that if you know at most six games were played, then because half these potential games were at each team's home field, you would correctly conclude that the probability that Team A won the series is exactly 1/2. If, however, you know that at most five games were played, then Team B, with its 3-to-2 home game advantage, is more likely to have won.

It follows that if *exactly* six games were played, Team A is more likely to have won!

Given that the answer to this puzzle does not depend on the degree of home-field advantage, you could arrive at this answer the following way. Assume that home-field advantage is overwhelming. Therefore, given that exactly six games were played, it is highly probable that there was only one upset—that is, only one game was won by the visiting team. Then that upset must have occurred in games 3, 4, 5 or 6, because otherwise the series would have ended at game 5. In only one of those cases (upset at game 6) would Team B have won the series, thus Team A has nearly a three-to-one advantage! This kind of "extremal reasoning" can be very useful in puzzle solving. As noted, however, it depends on the assumption that changing the parameter does not affect the solution.

2.6 Random Rice

This puzzle is an oldie but goodie. You'll need a bit of math background to be able to solve it in your head (or any other way, for that matter).

The volume of the ith squirt of rice is, by assumption, an independent uniformly random real number $\mathbf{X}_i$ between 0 and 1. We want to determine the expected value of the first j for which $\mathbf{X}_1 + \mathbf{X}_2 + \cdots + \mathbf{X}_j$ exceeds 1. That number will be at least 2, since the probability of getting a full cup on the first squirt is 0. But it might be 3 or even more if you are unlucky and begin with small squirts.

The key is to consider the fractional parts $\mathbf{Y}_1, \mathbf{Y}_2, \ldots$ of the partial sums. These numbers are also independent random numbers between 0 and 1, as you can easily see by noting that given the first i squirts, each possible value of $\mathbf{Y}_{i+1}$ arises from just one value of $\mathbf{X}_{i+1}$.

We can assume the machine's output is never exactly 0 or 1, since those events have probability 0 and hence do not affect the expected number of squirts. Then, as you squirt rice, the value of $\mathbf{Y}_i$ keeps going up until your rice total exceeds 1, at which point $\mathbf{Y}_i$ goes down. Thus, the probability that you'll need more than i squirts is exactly the probability that $\mathbf{Y}_1 < \mathbf{Y}_2 < \cdots < \mathbf{Y}_i$. This is just one of the $i!$ ways to order i numbers, so that probability is $1/i!$.

The expected value of any "counting" random variable is just the sum of the probabilities that that variable exceeds i, for all $i \geq 0$, so the expected value of the critical j is

$$\sum_{i=0}^{\infty} 1/i!,$$

which is the famous constant e. So on average, it takes exactly e (about 2.718281828459045) squirts to fill that rice cup.

2.7 Six with No Odds

This puzzle was communicated to me by MIT probabilist Elchanan Mossel, who came up with it as a problem for his undergraduate probability students before realizing that it was more subtle than he thought. The solution below is my own.

The first issue to be tackled is perhaps: Why is the answer not 3? Is this any different from rolling a die whose six faces are labeled with two 2s, two 4s, and two 6s? In that case, the answer would surely be 3, since the probability of "success" (rolling a 6) is 1/3, and when doing independent trials with probability p of success, the expected number of trials to reach success is $1/p$.

But it *is* different. When rolling an ordinary die until a 6 is obtained, conditioning on no odd numbers favors short experiments—if it took you a long time to roll your 6, you would probably have rolled some odd number on the way. Thus the answer to the puzzle should be less than 3.

It might help you to think about a series of experiments. If you repeatedly roll until a 6 appears, but ignore odd numbers, you will find that on the average it takes 3 (non-odd) rolls to get that 6. But the correct experiment is to *throw out the current series of rolls* if an odd number appears, then start a new series with roll number 1. Thus only series with no odds will count in your experiment.

That *gedankenexperiment* might give you an idea. Each series of rolls would end with either a 6 or an odd number. Does it matter? The final number in the sequence (1, 3, 5, or 6) is independent of the length of the series. Thus, by *reciprocity of independence* if you like, the length of the series is independent of what number caused it to end.

If you simply roll a die until either a 1, 3, 5, or 6 appears, it takes 3/2 rolls on average (since here the probability of "success" is 2/3). Since ending in a 6 has no effect on the number of trials, the answer to the original puzzle is that same 3/2.

2.8 Getting the Benz

You have only your author to blame for this last puzzle.

Your first thought, perhaps, is that since you have seen one of each number, all is fair and the odds have not changed.

But wait a second. The 1 you saw might have been yours, but the 2 and 3 were definitely submitted by your siblings. If you had guessed 2 or 3 instead of 1, you would now be plumb out of luck (for this round, anyway). So you did well to guess 1, thus your chances must be higher now.

Thinking of it another way, what is written on the two slips you did not see? There are $3^2 = 9$ possibilities, of which two—$(2, 3)$ and $(3, 2)$—get you the car, and three—(1,1), (2,2), and (3,3)—get you a rematch. You do not need a calculator to work out that this gives you quite a bit better than your original $1/5$ probability of getting the Benz.

Wonderful! You can already imagine your friends' jealous stares when you show up in your classic coupé. But something's nagging at your brain. It seems that your siblings, whatever they wrote on their slips, could all reason this way with the same peek that you got. Since you are all going for the same car, how can it be that you are all happy?

Acknowledgment

Research supported by NSF grant DMS–1600116.

References

[1] D. Hess. *Golf on the Moon*. Dover, Mineola, NY, 2014.
[2] C. F. Mosteller. *50 Challenging Problems in Probability*. Addison-Wesley, Reading, MA, 1965.

Index