Reading: How to win the lottery with geometry

This reading is an excerpt from Jordan Ellenberg's book *How Not to Be Wrong: The Power of Mathematical Thinking.* In this chapter, Ellenberg is explaining the true story of how two groups of people made money (yes, profit!) by playing the Michigan state lottery, a game called "Cash WinFall". Most of the ideas involved are basic probability theory – the particular design of this lottery game meant that on certain days, buying a ticket had *positive expected value.* So if you bought a lot of tickets on those days, you would very likely make money.

The people involved are :

a) A retired man from Michigan named Gerald Selbee, leading a group of "investors" who all pitched in money to buy high volumes of tickets

and

b) A group of undergraduate students from MIT who called themselves the "Random Strategies", led by a senior named James Harvey.

To play Cash WinFall, you pick 6 numbers between 1 and 46. The lottery then picks 6 winning numbers, and you get prizes based on how many of your numbers match the winners. (Ellenberg will later discuss a simpler version of this game, where you are supposed to pick 3 numbers between 1 and 7).

The reading begins with a comparison of the strategies of these two groups. Selbee's group buys a random assortment of tickets with an essentially random assortment of numbers, while Harvey's group chooses all the numbers themselves....

The notion of utility helps make sense of a puzzling feature of the Cash WinFall story. When Gerald Selbee's betting group bought massive quantities of tickets, they used Quic Pic, letting the lottery's computers pick the numbers on their slips at random. Random Strategies, on the other hand, picked their numbers themselves; this meant they had to fill out hundreds of thousands of slips by hand, then feed them through the machines at their chosen convenience stores one by one, a massive and incredibly dull undertaking.

The winning numbers are completely random, so every lottery ticket has the same expected value; Selbee's 100,000 Quic Pics would bring in the same amount of prize money, on average, as Harvey and Lu's 100,000 artisanally marked tickets. As far as expected value is concerned, Random Strategies did a lot of painful work for no reward. Why?

Consider this case, which is simpler but of the same nature. Would you rather have \$50,000, or would you rather have a 50/50 bet between losing \$100,000 and gaining \$200,000? The expected value of the bet is

$$(1/2) \times (-\$100{,}000) + (1/2) \times (\$200{,}000) = \$50{,}000,$$

the same as the cash. And there is indeed some reason to feel indifferent between the two choices; if you made that bet time after time after time, you'd almost certainly make \$200,000 about half the time and lose \$100,000 the other half. Imagine you alternated winning and losing: after two bets you've won \$200,000 and lost \$100,000 for a net gain of \$100,000, after four bets you're up \$200,000, after six bets \$300,000, and so on: a profit of \$50,000 per bet on average, just the same as if you'd gone the safe route.

But now pretend for a moment that you're not a character in a word problem in an economics textbook, but rather an actual person—an actual person who does not have \$100,000 cash on hand. When you lose that first bet and your bookie—let us say your big, angry, bald, power-lifting bookie—comes to collect, do you say, "An expected value calculation shows that it's very likely I'll be able to pay you back in the long run"? You do not. That argument, while mathematically sound, will not achieve its goals.

If you're an actual person, you should take the \$50,000.

What does all this have to do with Cash WinFall? As we said at the top, the expected dollar value of 100,000 lottery tickets is what it is, no matter which tickets you buy. But the variance is a different story. Suppose, for instance, I decide to go into the high-volume betting game, but I take a different approach; I buy 100,000 copies of the same ticket.

If that ticket happens to match 4 out of the 6 numbers in the lottery drawing, then I'm the lucky holder of 100,000 pick-4 winners, and I'm basically going to sweep up the entire \$1.4 million prize pool, for a tidy 600% profit. But if my set of numbers is a loser, I lose my whole \$200,000 pile. That's a high-variance bet, with a big chance of a big loss and a small chance of an even bigger win.

So "don't put all your money on one number" is pretty good advice—much better to spread your bets around. But wasn't that exactly what Selbee's gang was doing by using the Quic Pic machine, which chooses numbers at random?

Not quite. First of all, while Selbee wasn't putting all his money on one ticket, he *was* buying the same ticket multiple times. At first, that seems strange. At his most active, he was buying 300,000 tickets per drawing, letting the computer pick his numbers randomly from almost 10 million choices. So his purchases amounted to a mere 3% of the possible tickets; what are the odds he'd buy the same ticket twice?

Actually, they're really, really good. Old chestnut: bet the guests at a party that two people in the room have the same birthday. It had better be a good-sized party—say there are thirty people there. Thirty birthdays out of 365 options* aren't very many, so you might think it pretty unlikely that two of those birthdays would land on the same day. But the relevant quantity isn't the number of people: it's the number of *pairs* of people. It's not hard to check that there are 435 pairs of people,† and each pair has a 1 in 365 chance of sharing a birthday; so in a party that size you'd expect to see a pair sharing a birthday, or maybe even two

* 366 if you count leap days, but we're not going for precision here.

† The first person in the pair can be any of the 30 people, and the second any of the 29 who remain, giving 30 × 29 choices; but this counts each pair twice, since it counts {Ernie, Bert} and {Bert, Ernie} separately; so the right number of pairs is (30 × 29)/2 = 435.

pairs. In fact, the chance that two people out of thirty share a birthday turns out to be a little over 70%—pretty good odds. And if you buy 300,000 randomly chosen lottery tickets out of 10 million options, the chance of buying the same ticket twice is so close to 1 that I'd rather just say "it's a certainty" than figure out how many more 9s I'd need after "99.9%" to specify the probability on the nose.

And it's not just repeated tickets that cause the trouble. As always, it can be easier to see what's going on with the math if we make the numbers small enough that we can draw pictures. So let's posit a lottery draw with just seven balls, of which the state picks three as the jackpot combination. There are thirty-five possible jackpot combos, corresponding to the thirty-five different ways that three numbers can be chosen from the set 1, 2, 3, 4, 5, 6, 7. (Mathematicians like to say, for short, "7 choose 3 is 35.") Here they are, in numerical order:

123 124 125 126 127
134 135 136 137
145 146 147
156 157
167
234 235 236 237
245 246 247
256 257
267
345 346 347
356 357
367
456 457
467
567

Say Gerald Selbee goes to the store and uses the Quic Pic to buy seven tickets at random. His chance of winning the jackpot remains pretty small. But in this lottery, you also get a prize for hitting two out of three numbers. (This particular lottery structure is sometimes called the *Transylvanian lottery*, though I could find no evidence that such a game has ever been played in Transylvania, or by vampires.)

Two out of three is a pretty easy win. So I don't have to keep typing "two out of three," let's call a ticket that wins this lesser prize a *deuce*. If the jackpot drawing is 1, 4, and 7, for example, the four tickets with a 1, a 4, and some number *other* than 7 are all deuces. And besides those four, there are the four tickets that hit 1–7 and the four that hit 4–7. So twelve out of thirty-five, just over a third of the possible tickets, are deuces. Which suggests there are probably at least a couple of deuces among Gerald Selbee's seven tickets. To be precise, you can compute that Selbee has

5.3% chance of no deuces
19.3% chance of exactly one deuce
30.3% chance of two deuces
26.3% chance of three deuces
13.7% chance of four deuces
4.3% chance of five deuces
0.7% chance of six deuces
0.1% chance of all seven tickets being deuces.

The expected number of deuces is thus

$$5.3\% \times 0 + 19.3\% \times 1 + 30.3\% \times 2 + 26.3\% \times 3 + 13.7\% \times 4 + 4.3\% \times 5 + 0.7\% \times 6 + 0.1\% \times 7 = 2.4$$

The Transylvanian version of James Harvey, on the other hand, doesn't use the Quic Pic; he fills out his seven tickets by hand, and here they are:

124
135
167
257
347
236
456

Suppose the lottery draws 1, 3, and 7. Then Harvey's holding three deuces: 135, 167, and 347. What if the lottery draws 3, 5, 6? Then Harvey once again has three deuces among his tickets, with 135, 236, and 456. Keep trying possible combinations and you'll quickly see that Harvey's choices have a very special property: either he wins the jackpot, or he wins *exactly* three deuces. The chance that the jackpot is one of Harvey's seven tickets is 7 out of 35, or 20%. So he has a

20% chance of no deuces
80% chance of three deuces.

His expected number of deuces is

$$20\% \times 0 + 80\% \times 3 = 2.4$$

the same as Selbee's, as it must be. But the variance is much smaller; Harvey has very little uncertainty about how many deuces he's going to get. That makes Harvey's portfolio a lot more attractive to potential cartel members. Note especially: whenever Harvey doesn't get three deuces, he wins the jackpot. That means that Harvey's strategy *guarantees* a substantial minimum payoff, something the Quic-Pickers like Selbee can never do. Picking the numbers yourself can get rid of your risk while maintaining the reward—if you pick the numbers right.

And how do you do that? That is—literally, for once!—the million-dollar question.

First try: just ask your computer to do it. Harvey and his team were MIT students, presumably able to knock off a few dozen lines of code before their morning coffee. Why not just write a program to run through all combinations of 300,000 WinFall tickets to see which one provided the lowest-variance strategy?

That wouldn't be a hard program to write. The one small problem would be the way all matter and energy in the universe decayed into heat death by the time your program had handled the first tiny fragment of a microsliver of the data it was trying to analyze. From the point of view of a modern computer, 300,000 is not a very large number. But the objects that the proposed program has to pick through are not the 300,000 tickets—they are the possible sets of 300,000 tickets to be purchased from the 10 million possible Cash WinFall tickets. How many of those sets are there? More than 300,000. More than the number of subatomic particles that exist or have ever existed. A lot more. You've probably never even *heard* of a number as big as the number of ways to select your 300,000 tickets.*

What we're up against here is the dreaded phenomenon known by computer-science types as "the combinatorial explosion." Put simply: very simple operations can change manageably large numbers into absolutely impossible ones. If you want to know which of the fifty states is the most advantageous place to site your business, that's easy; you just have to compare fifty different things. But if you want to know which *route* through the fifty states is the most efficient—the so-called traveling salesman problem—the combinatorial explosion goes off, and you face difficulty on a totally different scale. There are about 30 vigintillion routes to choose from. In more familiar terms, that's 30 thousand trillion trillion trillion trillion trillion.

Boom!

So there'd better be another way to choose our lottery tickets to tamp down variance. Would you believe me if I told you it all came down to plane geometry?

WHERE THE TRAIN TRACKS MEET

Parallel lines don't meet. That's what makes them parallel.

But parallel lines sometimes *appear* to meet—think of a pair of train tracks, alone in an empty landscape, the two rails seeming to converge as your eyes follow them closer and closer to the horizon. (I find it helps to have some country music playing if you want a really vivid mental image here.) This is the phenomenon of *perspective*; when you try to depict the three-dimensional world on your two-dimensional field of vision, something has to give.

The people who first figured out what was going on here were the

* Unless you've heard of a googolplex. Now *that* is a big number, boy howdy.

people who needed to understand both how things are and how things look, and the difference between the two: namely, painters. The moment, early in the Italian Renaissance, at which painters understood perspective was the moment visual representation changed forever, the moment when European paintings stopped looking like your kid's drawings on the refrigerator door (if your kid mostly drew Jesus dead on the cross) and started looking like the things they were paintings of.*

How exactly Florentine artists like Filippo Brunelleschi came to develop the modern theory of perspective has occasioned a hundred quarrels among art historians, into which we won't enter here. What we know for sure is that the breakthrough joined aesthetic concerns with new ideas from mathematics and optics. A central point was the understanding that the images we see are produced by rays of light that bounce off objects and subsequently strike our eye. This sounds obvious to a modern ear, but believe me, it wasn't obvious then. Many of the ancient scientists, most famously Plato, argued that vision must involve a kind of fire that emanated from the eye. This view goes at least as far back as Alcmaeon of Croton, one of the Pythagorean weirdos we met in chapter 2. The eye must generate light, Alcmaeon argued: what other source could there be for the *phosphene*, the stars you see when you shut your eyes and press down on your eyeball? The theory of vision by reflected rays was worked out in great detail by the eleventh-century Cairene mathematician Abu 'Ali al-Hasan ibn al-Haytham (but let's call him Alhazen, as most Western writers do). His treatise on optics, the *Kitab al-Manazir*, was translated into Latin and taken up eagerly by philosophers and artists seeking a more systematic understanding of the relation between sight and the thing seen. The main point is this: a point P on your canvas represents a *line* in three-dimensional space. Thanks to Euclid, we know there's a unique line containing any two specified points. In this case, the line is the one containing P and your eye. Any object in the world that lies on that line gets painted at point P.

Now imagine you're Filippo Brunelleschi standing out on the flat

* Or at least they looked like certain kinds of optical representations of the things they were paintings of, which over the years we've come to think of as realistic; what counts as "realism" has been the subject of hot contention among art critics for about as long as there's been art criticism.

prairie, the canvas on an easel in front of you, painting the train tracks.* The track consists of two rails, which we call R_1 and R_2. Each one of these rails, drawn on the canvas, is going to look like a line. And just as a point on the canvas corresponds to a line in space, a line on the canvas corresponds to a plane. The plane P_1 corresponding to R_1 is the one swept out by the lines joining each point on the rail to your eye. In other words, it's the unique plane containing both your eye and the rail R_1. Similarly, the plane P_2 corresponding to R_2 is the one containing your eye and R_2. Each of the two planes cuts the canvas in a line, and we call these lines L_1 and L_2.

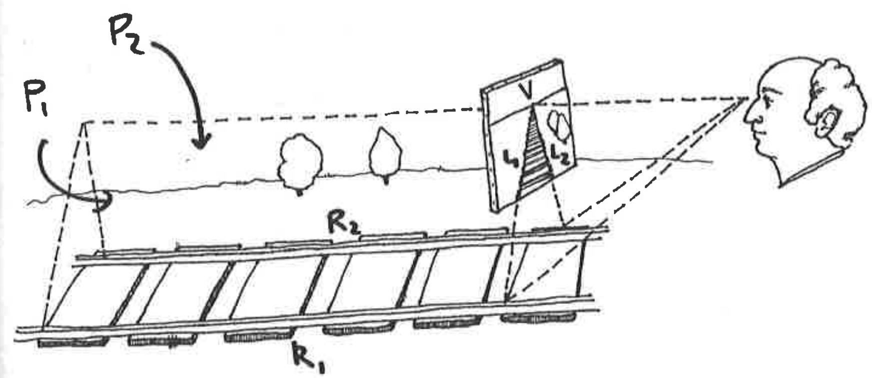

The two rails are parallel. *But the two planes are not.* How could they be? They meet at your eye, and parallel planes do not meet anywhere. But planes that aren't parallel have to intersect in a line. In this case, the line is horizontal, emanating from your eye and proceeding parallel to the train tracks. The line, being horizontal, does not meet the prairie—it shoots out toward the horizon, never touching the ground. But—and here is the point—it meets the canvas, at some point V. Since V is on the plane P_1, it must be on the line L_1 where P_1 cuts the canvas. And since V is also on P_2, it must be on L_2. In other words, V is the point on the canvas where the painted train tracks meet. In fact, any straight path on the prairie that runs parallel to the train tracks will look, on the canvas, like

* Anachronistic, okay, but just go with it.

a line through V. V is the so-called vanishing point, the point through which the paintings of all lines parallel to the tracks must pass. In fact, every pair of parallel tracks determines some vanishing point on the canvas; where the vanishing point is depends on which direction the parallel lines are going. (The only exceptions are pairs of lines parallel to the canvas itself, like the slats between the rails—they'll still look parallel in your painting.)

The conceptual shift that Brunelleschi made here is the heart of what mathematicians call projective geometry. Instead of points in the landscape, we think of lines through our eye. At first glance, the distinction might seem purely semantic; each point on the ground determines one and only one line between the point and our eye, so what does it matter whether we think about the point or think about the line? The difference is just this: there are more lines through our eye than there are points on the ground, because there are *horizontal* lines, which don't intersect the ground at all. These correspond to the vanishing points on our canvas, the places where train tracks meet. You might think of this line as a point on the ground that is "infinitely far away" in the direction of the tracks. And indeed, mathematicians usually call them *points at infinity*. When you take the plane Euclid knew and glue on the points at infinity, you get the *projective plane*. Here's a picture of it:

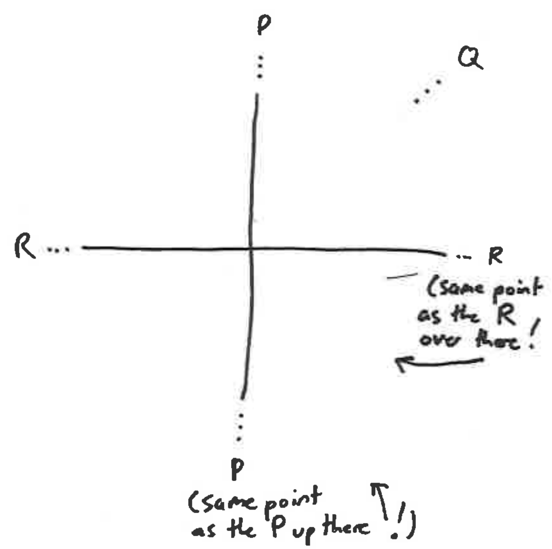

Most of the projective plane looks just like the regular flat plane you're used to. But the projective plane has more points, those so-called points at infinity: one for each possible direction along which a line can be oriented in the plane. You should think of the point P, which corresponds to the vertical direction, as being infinitely high up along the vertical axis—but also infinitely *low down* along the vertical axis. In the projective plane, the two ends of the y-axis *meet* at the point at infinity, and the axis is revealed to be not really a line but a circle. In the same way, Q is the point that's infinitely far northeast (or southwest!) and R is the point at the end of the horizontal axis. Or rather, at *both* ends. If you travel infinitely far to the right, until you arrive at R, and then keep on going, you find yourself still traveling rightward but now heading back toward the center from the left edge of the picture.

This kind of leaving-one-way-and-coming-back-the-other enthralled the young Winston Churchill, who recalled vividly the one mathematical epiphany of his life:

> I had a feeling once about Mathematics, that I saw it all—Depth beyond depth was revealed to me—the Byss and the Abyss. I saw, as one might see the transit of Venus—or even the Lord Mayor's Show, a quantity passing through infinity and changing its sign from plus to minus. I saw exactly how it happened and why the tergiversation was inevitable: and how the one step involved all the others. It was like politics. But it was after dinner and I let it go!

In fact, point R is not just the endpoint of the horizontal axis, but of *any* horizontal line. If two different lines are both horizontal, they are parallel; and yet, in projective geometry, they meet, at the point at infinity. David Foster Wallace was asked in a 1996 interview about the ending of *Infinite Jest*, which many people found abrupt: Did he, the interviewer asked, avoid writing an ending because he "just got tired of writing it"? Wallace replied, rather testily: "There is an ending as far as I'm concerned. Certain kinds of parallel lines are supposed to start converging in such a way that an 'end' can be projected by the reader somewhere beyond the right frame. If no such convergence or projection occurred to you, then the book's failed for you."

The projective plane has the defect that it's kind of hard to draw, but the advantage that it makes the rules of geometry much more agreeable. In Euclid's plane, two different points determine a single line, and two different lines determine a single intersection point—unless they're parallel, in which case they don't meet at all. In mathematics, we like rules, and we don't like exceptions. In the projective plane, you don't have to make any exceptions to the rule that two lines meet at a point, because parallel lines meet too. Any two vertical lines, for instance, meet at P, and any two lines pointing northeast to southwest meet at Q. Two points determine a single line, two lines meet at a single point, end of story.* It's perfectly symmetrical and elegant in a way that classical plane geometry is not. And it's not coincidence that projective geometry arose naturally from attempts to solve the practical problem of depicting the three-dimensional world on a flat canvas. Mathematical elegance and practical utility are close companions, as the history of science has shown again and again. Sometimes scientists discover the theory and leave it to mathematicians to figure out why it's elegant, and other times mathematicians develop an elegant theory and leave it to scientists to figure out what it's good for.

One thing the projective plane is good for is representational painting. Another is picking lottery numbers.

A TINY GEOMETRY

The geometry of the projective plane is governed by two axioms:

> Every pair of points is contained in exactly one common line.
> Every pair of lines contains exactly one common point.

Once mathematicians had found *one* kind of geometry that satisfied these two perfectly tuned axioms, it was natural to ask whether there

* But if the lines containing R are all horizontal, and the lines containing P are all vertical, what is the line through R and P? It is a line we haven't drawn, the *line at infinity*, which contains all the points at infinity and none of the points of the Euclidean plane.

were any more. It turns out there are a lot. Some are big, some are small. The very tiniest is called the *Fano plane*, after its creator, Gino Fano, who in the late nineteenth century was one of the first mathematicians to take seriously the idea of finite geometries. It looks like this:

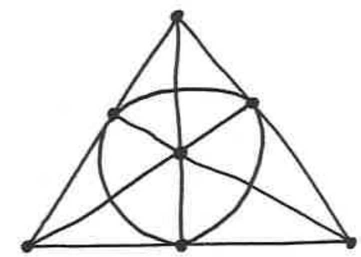

This is a small geometry indeed, consisting of only seven points! The "lines" in this geometry are the curves shown in the diagram; they're small, too, possessing only three points each. There are seven of them, six of which *look* like lines and the other of which looks like a circle. And yet this so-called geometry, exotic as it is, satisfies axioms 1 and 2 just as well as Brunelleschi's plane did.

Fano had an admirably modern approach—he had, to use Hardy's phrase, "the habit of definition," avoiding the unanswerable question of what geometry *really was*, and asking, instead: Which phenomena behave like geometry? In Fano's own words:

> A base del nostro studio noi mettiamo una *varietà* qualsiasi di enti di qualunque natura; enti che chiameremo, per brevità, punti indipendentemente però, ben inteso, dalla loro stessa natura.

That is:

> As a basis for our study we assume an arbitrary *collection* of entities of an arbitrary nature, entities which, for brevity, we shall call points, but this is quite independent of their nature.

For Fano and his intellectual heirs, it doesn't matter whether a line "looks like" a line, a circle, a mallard duck, or anything else—all that matters is that lines *obey the laws* of lines, set down by Euclid and his successors. If it walks like geometry, and it quacks like geometry, we call it geometry. To one way of thinking, this move constitutes a rupture between mathematics and reality, and is to be resisted. But that view is too conservative. The bold idea that we can think geometrically about systems that don't look like Euclidean space,* and even call these systems "geometries" with head held high, turned out to be critical to understanding the geometry of the relativistic space-time we live in; and nowadays we use generalized geometric ideas to map Internet landscapes, which are even further removed from anything Euclid would recognize. That's part of the glory of math; we develop a body of ideas, and once they're correct, *they're correct*, even when applied far, far outside the context in which they were first conceived.

For example: here's Fano's plane again, but with the points labeled by the numbers 1 through 7:

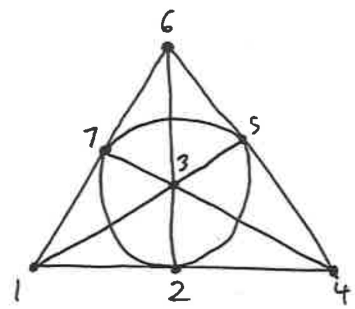

* To be fair, there is another sense in which the Fano plane really does look like more traditional geometry. Descartes taught us how to think of points on the plane as pairs of *coordinates* x and y, which are real numbers; if you use Descartes's construction but draw your coordinates from number systems other than the real numbers, you get other geometries. If you do Cartesian geometry using the Boolean number system beloved of computer scientists, which has only two numbers, the bits 0 and 1, you get the Fano plane. That's a beautiful story, but it's not the story we're telling just now. See the endnotes for a little more of it.

Look familiar? If we list the seven lines, recording for each the set of three points that constitute it, we get:

124
135
167
257
347
236
456

This is none other than the seven-ticket combo we saw in the last section, the one that hits each pair of numbers exactly once, guaranteeing a minimum payoff. At the time, that property seemed impressive and mystical. How could anyone have come up with such a perfectly arranged set of tickets?

But now I've opened the box and revealed the trick: it's simple geometry. Each pair of numbers appears on exactly one ticket, because each pair of points appears on exactly one line. It's just Euclid, even though we're speaking now of points and lines Euclid would not have recognized as such.

I'M SORRY, DID YOU SAY "BOFAB"?

The Fano plane tells you how to play the seven-number Transylvanian lottery without taking on any risk, but what about the Massachusetts lottery? There are lots of finite geometries with more than seven points, but none, unfortunately, that precisely meet the requirements of Cash WinFall. Something more general is needed. The answer doesn't come directly from Renaissance painting or Euclidean geometry, but from another unlikely source—the theory of digital signal processing.

Suppose I want to send an important message to a satellite, like "Turn on right thruster." Satellites don't speak English, so what I'm actually sending is a sequence of 1s and 0s, what computer scientists call *bits*: