Math Puzzles

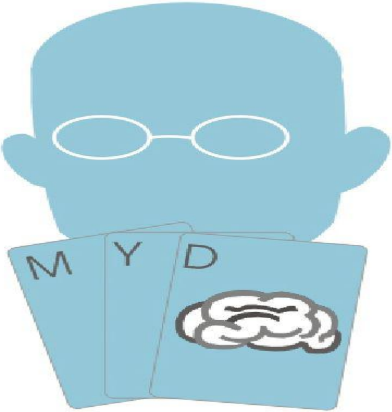

Classic riddles in counting, geometry, probability, and game theory

Presh Talwalkar

Math puzzles: classic riddles in counting,

geometry, probability, and game theory

Presh Talwalkar

Copyright Presh Talwalkar, 2012

All rights reserved

Acknowledgements

I want to thank readers of Mind Your Decisions for their continued support. I always loved math puzzles and I'm fortunate for the community of readers that have suggested puzzles to me and offered ingenious solutions.

These puzzles are a collection of problems that I have read over the years in books, math competitions, websites, and emails. I have credited original sources when aware, but if there are any omissions please let me know at presh@mindyourdecisions.com

Table of Contents

Section 1: Counting and geometry problems

The following 25 puzzles deal with classic riddles about counting and geometry.

Puzzle 1: Ants on a Triangle

Three ants are positioned on separate corners of a triangle.

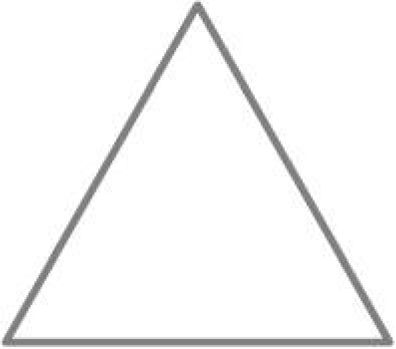

If each ant moves along an edge toward a randomly chosen corner, what is the chance that none of the ants collide?

How would the problem generalize if there were n ants positioned on the vertices of a regular n-gon, again each moving along an edge to a random adjacent vertex?

Answer to Puzzle 1: Ants on a Triangle

In order that none of the ants collide, they must all move in the same direction. That is, all of the ants must move in either clockwise or counter-clockwise towards a new corner.

This can be seen by inductive reasoning: whichever orientation ant 1 picks, ant 2 must pick the same orientation to avoid a collision, and then ant 3 must do the same thing as well.

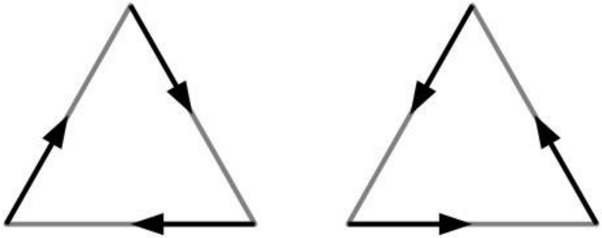

Therefore there are 2 different ways that the ants can avoid running into each other.

As each ant can travel in to 2 different directions, there are $2^3 = 8$ total possible ways the ants can move.

The probability the ants do not collide is 2/8 = 25%.

Extension to general n-gon

An interesting extension is to ask what would happen to 4 ants on the vertices of a quadrilateral? Or more generally, if there are n ants on an n-gon?

The general problem can be solved by the same logic.

Again, the ants can only avoid collision if they all move in the same orientation--either clockwise or counter-clockwise. So again there are only 2 safe routes the ants as a group can take.

The total number of routes the ants can take is also easy to count. Each ant can choose between 2 adjacent vertices, so there are 2^n possibilities ways the ants could choose to travel.

The probability that none of the ants will collide is $2/2^n = 1/2^{n-1}$

For example, on an 8-sided octagon, the probability that none of the ants will collide is a mere 1/128 = 0.0078125, which is less than one percent.

For larger polygons it will be rare that the ants do not collide, but not impossible.

Puzzle 2: Three brick problem

This is a very simple problem that tests a bit of your creative ability.

How can you measure the diagonal of a brick without using any formula, if you have three bricks and a ruler?

Answer to Puzzle 2: Three brick problem

There is a remarkably easy way to find the diagonal.

What you need to do is stack two bricks, one on top of each other, and then place the third brick next to the bottom brick.

The idea is you create an empty space in which a fourth brick could be placed. Here is a diagram to illustrate:

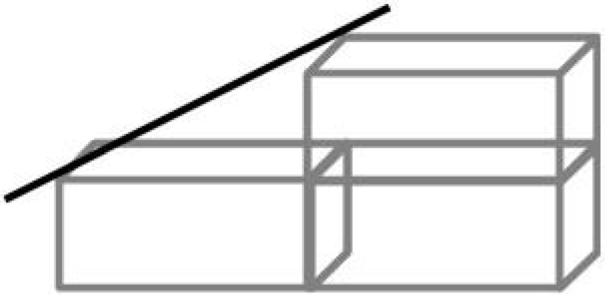

Now you can measure the length of the diagonal by measuring the length of the empty space using a ruler.

No Pythagorean theorem or geometry formulas required!

Puzzle 3: World's best tortilla problem

You start out with a round tortilla, as depicted below.

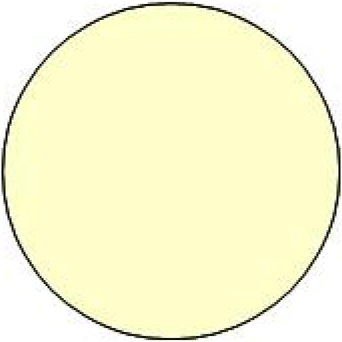

Your job is to divide the tortilla into eight equal pieces, using only cuts made in a straight line.

What is the minimum number of cuts you need to make?

(credit: problem is adapted from The World's Best Puzzles)

Answer to Puzzle 3: World's best tortilla problem

You only need one cut! The trick to this problem is you can fold the tortilla three times in half and then make one cut (fold along the dotted lines, and then cut along the dark line in the following diagram.

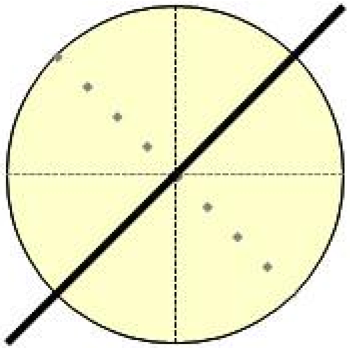

I came across this problem when I was making homemade baked tortilla chips. Most instructional cooking videos show people inefficiently making 4 cuts which I found somewhat annoying.

Puzzle 4: Slicing up a pie

Alice and Bob are preparing for a holiday party, and each has a pie to slice up into pieces.

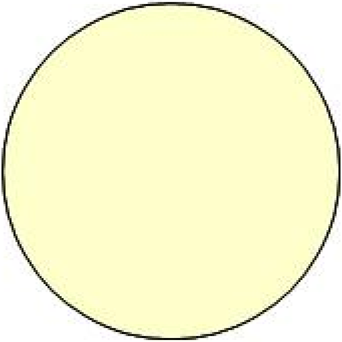

They decide to have a little contest to make things fun. Each person is allowed to make 3 cuts of the pie with a knife, and whoever ends up with more pieces is the winner. They agree stacking is not allowed, but that “center” pieces without the crust are permissible.

How many pieces can be made using 3 cuts? What about 4 cuts, or more generally n cuts?

Answer to Puzzle 4: Slicing a pie

When you make one cut, you can create 2 halves. With two cuts, you can slice through each half again, to make 4 pieces.

The third cut is a bit trickier. What you want is to cross the previous cuts without going through the intersection point. Every time you intersect a previous cut you create another section (piece) of the pie, as in the following diagram.

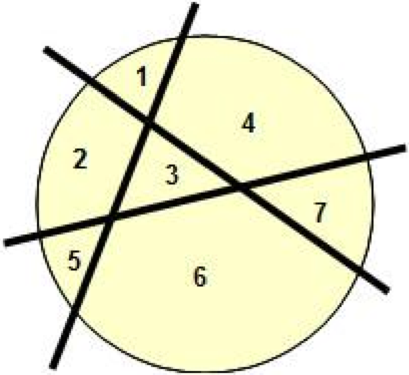

So with 3 cuts, you can make 7 pieces in all.

We can now generalize. The first cut makes the pie into 2 pieces. But after that, making the cut *n* will add on *n* new pieces to the pie. Thus, we know that on cut *n* the total number of pieces can be calculated by the formula:

$f(n) = 2 + 2 + 3 + 4 + \ldots + n = 1 + (1 + 2 + 3 + \ldots) = 1 + n(n + 1)/2$

This sequence has a special name. The number of cuts you can make with *n* cuts is known as lazy caterer's sequence.

Puzzle 5: Measuring ball bearings

This is a classic puzzle about weighing.

You are given a container that contains hundreds of ball bearings, amassing to exactly 24 ounces.

You have a balance but no weights for the scale.

You want to measure exactly 9 ounces. How can you do it?

Answer to Puzzle 5: Measuring ball bearings

If you could count the number of ball bearings, you could get a unit weight for 1 bearing and proceed by counting. But the ball bearings are too numerous, and you can figure it out quicker by using several weighings.

Here is one way to do it in five steps.

1. Divide the balls into two equal piles using the balance (12 ounces on each side)

2. Remove the ball bearings from the scale. Divide one of the 12 ounce piles into two equal piles using the scale (6 ounces on each side)

3. Set aside one of the 6 ounce piles

4. Divide the other 6 ounce pile into two piles (3 ounces on each side)

5. Combine a 3 ounce pile with a 6 ounce pile to get to 9 ounces

Puzzle 6: Paying an employee in gold

You have a solid gold bar, marked into 7 equal divisions as follows:

| – | – | – | – | – | – | – |

You need to pay an employee each day for one week. He asks to be paid exactly 1 piece of the gold bar per day.

The problem is you don't trust him enough to prepay him, and he would prefer not to be paid late.

If you can only make 2 cuts in the bar, can you figure out a way to make the cuts so your worker gets paid exactly one gold piece every day?

Answer to Puzzle 6: Paying an employee in gold

You can pay the employee if you cut the pieces in the correct spots. If the gold bar is labeled as follows:

| 1 | 2 | 3 | 4 | 5 | 6 | 7 |

Then you want to make the cuts between pieces 1 and 2, and between pieces 4 and 5. So now you have pieces:

|1| |2|3| |4|5|6|7|

Now you have a 1-block piece, a 2-block piece, and a 4-block piece. Here is how you can pay the employee one piece of gold for each day during the week:

Day 1: Give him the 1-block piece

Day 2: Trade him the 2-block piece for the 1-block piece

Day 3: Give him back the 1-block piece

Day 4: Trade him the 4-block piece for the 1 and 2-block pieces

Day 5: Give him the 1-block piece

Day 6: Trade him the 2-block piece for the 1-block piece

Day 7: Give him the 1-block piece back

As you can see, the worker will be paid 1 block each day.

Puzzle 7: Leaving work quickly

Alice and Bob were ready to leave the office when their mean boss assigned them more work.

The boss told them to do the following boring things before they could go home:

1. Manually copy pages from bound books
2. Audit numbers in a spreadsheet
3. Fax documents to another office

Each task takes 40 minutes to complete, and only one person can work on a task at a time (the office only has one copy machine, one fax machine, and auditing cannot be done simultaneously).

How quickly can they complete their work and go home?

Answer to Puzzle 7: Leaving work quickly

At first thought it seems like the tasks will require 80 minutes: in the first 40 minutes, each does one task, and in the last 40 minutes someone finishes the last task.

But let us diagram the chores using something called a Gantt chart which shows what is being done in each interval of time:

	Time	
	0 – 40	40 – 80
Alice	Audit	Copy
Bob	Fax	

You will notice in the second 40 minutes that only one person is working while the other is doing nothing. This chart should give us an idea of how to work more efficiently: both people need to be working simultaneously the entire time.

So imagine they split up the tasks into 20 minute intervals, and divide the tasks as follows:

	Time		
	0 – 20	20 – 40	40 – 60
Alice	Audit	Audit	Fax
Bob	Fax	Copy	Copy

Very curiously by splitting up one of the tasks (in this case, faxing), they are able to finish all of the work in 60 minutes!

In fact this really this should not be too big of a surprise: there are 3 chores that take 40 minutes for a total of 120 minutes. With two people working it should take no more than 60 minutes.

This type of problem is based on an old math puzzle about cooking three hamburgers on two grills: see details on page 133 and 134 of this pdf. I never liked this problem as much because you can't split up the task of grilling a burger: if you start cooking something and let it rest, it will keep cooking even if it is not on the flame. Nevertheless, the mathematical principle is useful for other problems.

Puzzle 8: Science experiment

A chemistry teacher offers his class an experiment for extra credit. To complete the lab, students are to keep bacteria in a special chamber for exactly 9 minutes.

The sadistic part is the teacher only gives the students a 4-minute and a 7-minute hourglass with which to measure time. There are no other time-measuring instruments, as wristwatches and cell phones are confiscated.

To complete the lab, the bacteria can be stored in small intervals of time, but the total time that it should be in the chamber must be 9 minutes.

Extra credit will only be awarded to the student or students that complete the lab first.

What is the shortest time the experiment can be completed?

A) 9 minutes
B) 12 minutes
C) 18 minutes
D) 21 minutes

Answer to Puzzle 8: Science experiment

The multiple choice of answers is a bit of a distractor. The experiment can be done in 9 minutes as the hourglasses can be used to measure this amount of time--see chart below.

Time	Action
0	Turn over both hourglasses
4	Turn over 4 min hourglass
7	Turn over 7 min hourglass
8	Turn over 7 min hourglass (again)
9	Take out the sample

The practical issue is how quickly students will realize the solution.

Puzzle 9: Elevator malfunctioning

An elevator in my office building of 65 floors is malfunctioning.

Whenever someone wants to go up, the elevator moves up by 8 floors if it can. If the elevator cannot move up by 8 floors, it stays in the same spot (if you are on floor 63 and press up, the elevator stays on floor 63).

And whenever someone wants to go down, the elevator moves down by 11 floors if it can. If it cannot, then the elevator stays in the same spot. (if you press down from floor 9, the elevator stays on floor 9).

The elevator starts on floor 1. Is it possible to reach every floor in the building?

How many times would you have to stop to reach the 60th floor, if you started on floor 1?

Answer to Puzzle 9: Elevator malfunctioning

It is possible to reach every floor in the building.

I used a spreadsheet to illustrate exactly how this is possible.

On the table below, every floor can be reached by a combination of moving up by 8 floors and moving down by 11 floors. (For simplicity, we can imagine the elevator has an "UP" button and a "DOWN" button).

The horizontal rows show an elevator moving up by 8 floors at a time, and the vertical columns are when the elevator moves down by 11 floors at a time.

	Pushing the “UP” button: go up by 8 floors																	
Pushing the “DOWN” button: go down by 11 floors	1	9	17	25	33	41	49	57	65									
			6	14	22	30	38	46	54	62								
				3	11	19	27	35	43	51	59							
						8	16	24	32	40	48	56	64					
							5	13	21	29	37	45	53	61				
								2	10	18	26	34	42	50	58			
										7	15	23	31	39	47	55	63	
											4	12	20	28	36	44	52	60

You can see that every floor is attainable.

The harder part is to see why this actually works.

The reason has to do with number theory. We are essentially looking for integers solutions to the following equation:

$$8x - 11y = \text{floor number}$$

These types of equations are known as linear Diophantine equations. For a general equation, $ax+by = c$, solutions exist if and only if c is a multiple of the greatest common divisor of a and b.

In our case, 8 and 11 are relatively prime and have a greatest common divisor of 1, thus there are infinitely many solutions to the equation. The tricky part is verifying that every floor can be reached without going above floor 65 or going below floor 1, which is what the table above demonstrates.

How many times would you have to stop to reach the 60th floor, if you started on floor 1?

I got the idea for this question when I noticed floor 60 is the farthest away from floor 1 in the table.

I calculated the quickest route to get to floor 60 is to take 24 stops along the way: go across the first row, then move down in a step ladder fashion until you reach the bottom row of the table, and then move across the last row.

The floor sequence is: 1, 9, 17, 25, 33, 41, 49, 57, 65, 54, 43, 32, 21, 10, 18, 7, 15, 4, 12, 20, 28, 36, 44, 52, and finally 60.

That's a long way to ride to get to your floor–that floor better hope someone fixes the elevator quickly!

(Credit: this puzzle is adapted from here)

Puzzle 10: Ants and honey

The shortest distance between two points on a plane is a straight line. But finding the shortest distance on other surfaces is a more interesting problem.

Here is a puzzle that is harder than it sounds.

In a rectangular box, with length 30 inches and height and width 12 inches, an ant is located on the middle of one side 1 inch from the bottom of the box.

There is a drop of honey at the opposite side of the box, on the middle of one side, 1 inch from the top.

Here is a picture that illustrates the position of the ant and the honey.

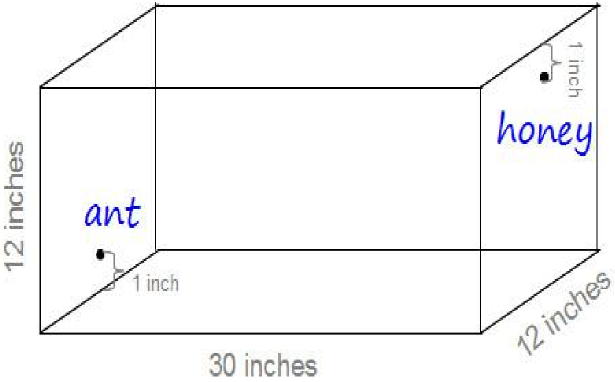

Let's say the ant is hungry and wants food quickly.

What is the shortest distance the ant would need to crawl to get the honey?

Answer to Puzzle 10: Ants and honey

If the ant crawls 1 inch down, then 30 inches across the bottom, then 11 inches up, it will travel 42 inches. But this is not the shortest distance.

The solution is found by unfolding the box and then finding the shortest path between the ant and the honey.

There are actually 4 ways to "flatten" the box (shown here). But only one method corresponds to the shortest distance as follows:

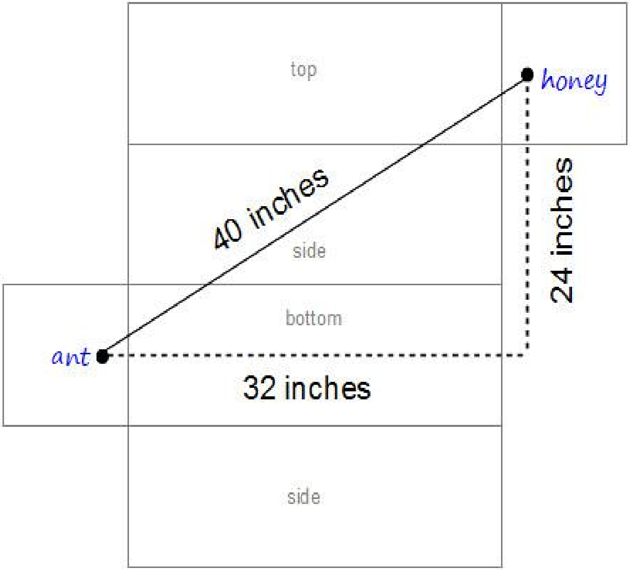

The distance between the points can be found using the Pythagorean theorem. For a triangle with legs 32 and 24, the hypotenuse--and shortest distance--is 40 inches.

Puzzle 11: Camel and bananas

This is a classic puzzle that I really enjoy.

You want to transport 3,000 bananas across 1,000 kilometers. You have a camel that can carry 1,000 bananas at most. However, the camel must eat 1 banana for each kilometer that it walks.

What is the largest number of bananas that can be transported?

Answer to Puzzle 11: Camel and bananas

The camel cannot carry all the bananas as it would eat them all in transport. Therefore, the bananas must be transported in shifts.

With 3,000 bananas, the camel will need to double back two times to carry the three different heaps of 1,000 bananas.

To carry the initial heap by 1 kilometer, the camel will need to make 5 trips and eat 5 bananas as follows:

> --Carry 1,000 bananas by 1 kilometer (eats 1 banana)
> --Return 1 kilometer to the beginning (eats 1 banana)
> --Carry the next 1,000 bananas by 1 kilometer (eats 1 banana)

--Return again 1 kilometer to the beginning (eats 1 banana)
--Carry the remaining bananas by 1 kilometer (eats 1 banana)

Notice that after moving 1 kilometer, the camel has eaten 5 of the bananas.

This process can be repeated and the camel will slowly transport and eat the bananas at the rate of 5 bananas per kilometer.

But after 200 kilometers, something important happens. At this point, the camel will have eaten 200 x 5 = 1,000 bananas, leaving just 2,000 remaining.

Because the camel can carry 1,000 bananas at a time, the camel will only need to double back once. To carry the remaining heap by 1 kilometer, the camel will only need to eat 3 bananas as follows:

--Carry 1,000 bananas by 1 kilometer (eats 1 banana)
--Return 1 kilometer to the beginning (eats 1 banana)
--Carry the remaining bananas by 1 kilometer (eats 1 banana)

The second leg therefore requires just 3 bananas per kilometer.

How long will this be necessary? Notice that after 333 1/3 kilometers, the camel has devoured another 1,000 bananas.

At this point, there are just 1,000 bananas left: the camel can make the remaining journey without doubling back. This means the camel can carry all the remaining 1,000 bananas and complete the trip.

How much of the trip remains? The camel went 200 kilometers and then 333 1/3 kilometers, so there are 466 2/3 kilometers remaining.

Thus, the camel will devour 466 2/3 bananas to complete the journey, meaning 533 1/3 bananas can survive and be transported.

Here is a visual representation of the journey:

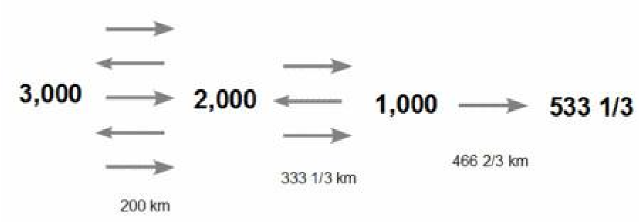

(credit: graphic inspired by this graphic)

Puzzle 12: Coin tossing carnival game

In one carnival game, you are to toss a coin on a table top marked with a grid of squares. You win if the coin lands without touching any lines–that is, the coin lands entirely inside one of the squares, as pictured below.

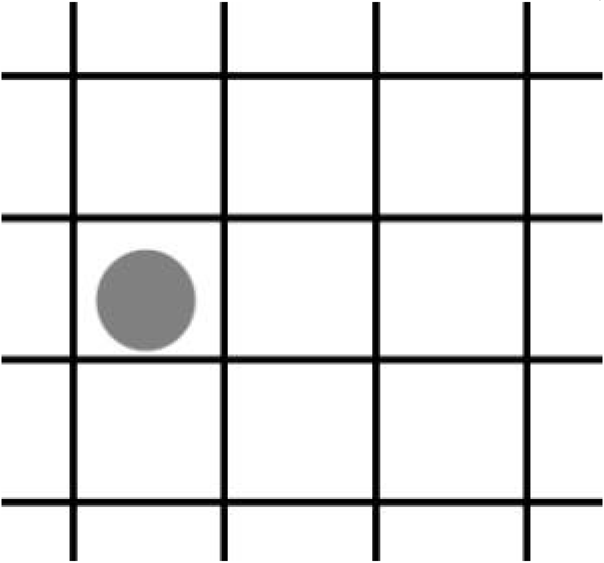

If the squares measure 1.5 inches per side, and the coin has a diameter of 1 inch, what is the chance you will win? Assume you can always get the coin to land somewhere on the table.

Extension: find a formula if the squares measure S inches per side and the coin measures D inches in diameter.

Answer to Puzzle 12: Carnival coin tossing game

The correct answer for this game is 1/9.

Let us solve the general case to see why. For the coin not to intersect any part of the grid, it must be the case that the circle's center is located sufficiently far enough away from the grid lines. These are all winnable points.

We can find the area of the winnable points and divide that by the total area of a square from the grid to calculate the probability of winning.

Here is a diagram that can help.

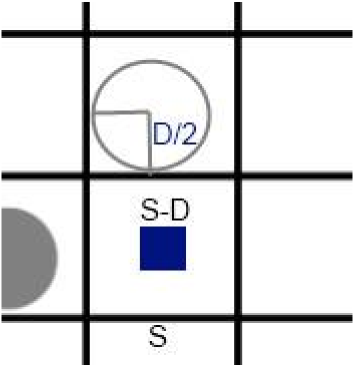

The winning points are the square with side S – D. This is found because the circle's center must be more than D/2 distance from all sides of the edge of a gridline. Hence we see the circle's center must lie in a square with side S - 2 (D/2) = S – D.

The area of the square for winning points is $(S – D)^2$. The area of a square for a gridline is S^2

The probability of winning is the ratio of these areas, which is $[(S\text{-}D)/S]^2$.

For a square of 1.5 inches, and a circle of diameter 1 inch, we find the probability of winning is $((0.5)/1.5)^{2} = 1/9$

Puzzle 13: Rope around Earth puzzle

This is a fun problem that first appeared in a 1702 book written by the philosopher William Whiston.

This problem is about two really, really long ropes A and B.

Rope A is long enough that it could wrap around the Earth's equator and fit snugly, like a belt (let's say 25,000 miles).

Rope B is just a bit longer than rope A. Rope B could wrap around the Earth equator from 1 foot off the ground.

How much longer is rope B than A?

(assume the earth is a perfect sphere)

Extension: let's say that rope C can wrap around an equatorial line for a sphere that's as big as the planet Jupiter (about 273,000 miles). Rope D is just a bit longer, and it can do the same thing from 1 foot off the ground.

How much longer is rope D than C?

Answer to Puzzle 13: Rope around Earth puzzle

The surprising part is that both questions have the same answer!

To see why, suppose that r is the radius of the Earth. Then, according to the setup, the larger rope B would have a radius of $r + 1$.

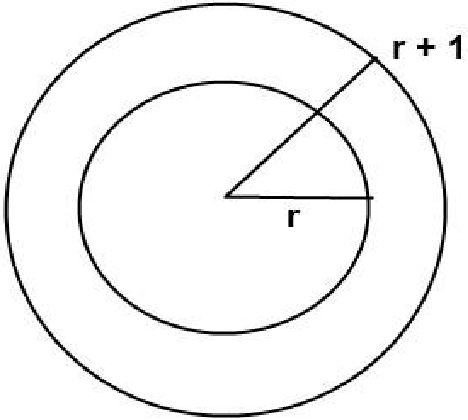

We can calculate how much longer rope *B* is by subtracting the difference of the circumferences of the two circles. The larger rope has circumference 2 π (*r*+1) and the smaller rope has one of 2 π *r*

$$2\pi(r+1) - 2\pi r = 2\pi = \text{about } 6.28 \text{ feet}$$

Therefore, rope *B* is longer by 6.28 feet.

But notice the remarkable thing: the answer does not depend on the radius of the circle! This means we have solved the problem for any size sphere (or one might say for every planet or spherically shaped object).

Hence, for Jupiter, rope *D* is also longer than *C* by about 6.28 feet.

Puzzle 14: Dividing a rectangular piece of land

A father is splitting up land among his two sons in estate planning. How can he divide the land fairly?

One approach is to split the land evenly. But even this method can get complicated if we add some realistic assumptions. This puzzle illustrates why splitting land can be a mind-boggling exercise.

Suppose your father owns a rectangular piece of land, but the city has bought a small rectangular patch of it for its public use.

You and your brother are to split up the land equally using only a single straight line to divide the area. How can this be done?

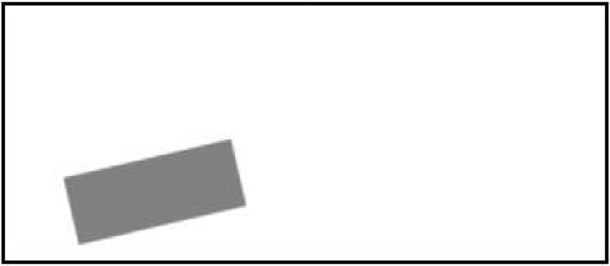

Trivia

As a bit of history, this puzzle is sometimes used as an interview brain teaser or technical question when testing job seekers on their problem solving ability.

It is sometimes stated in the following terms: how can you split in half a rectangular piece of cake, with a small rectangular piece removed, using a single cut from a knife?

Answer to Puzzle 14: Dividing a rectangular piece of land

The elegant mathematical solution requires a small trick about geometry. The trick is that any line passing through the center of a rectangle bisects its area.

(A line through the center of a rectangle either creates two equal triangles--if it is a diagonal--or it creates two equal trapezoids or rectangles)

The original rectangular plot of land has infinitely many lines passing through the center that bisect its area. But once you remove a small rectangular plot, there is only one line that bisects the area--namely, the line that passes through the centers of both rectangles. This line bisects both the original plot and the removed rectangular plot, and consequently splits the land evenly.

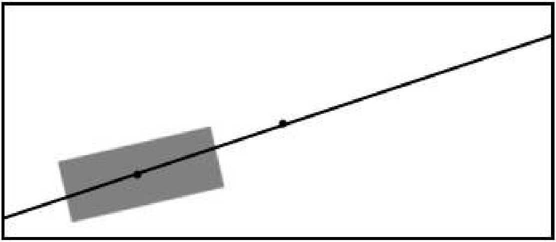

Another creative way was thought up by one of the Mind Your Decisions readers. Joe explained how he solved the puzzle on the spot.

I was asked this question during a recent job interview. My way of coming up with a solution was to rephrase the original puzzle by replacing rectangles with circles - i.e., a circle within a circle. When looking at the puzzle in this way, it's more intuitive to see a line connecting the two centers being the best answer, and then you can extend the analogy to the rectangles.

Now that's definitely what I call out of the box thinking.

Puzzle 15: Dividing land between four sons

This is one of my all-time favorite puzzles. Give it an honest effort before reading the answer.

A father dies and wants to divide his land evenly amongst four sons. The plot of land has the following unusual shape:

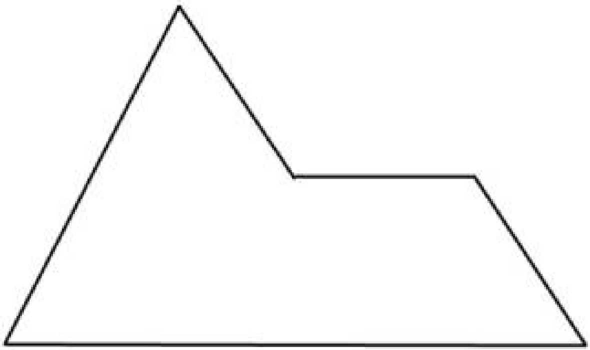

How can you divide the land into four equal parts, using only straight lines?

Answer to Puzzle 15: Dividing land between four sons

I came across this puzzle when it was presented to gifted math students.

Several of the high school students then were able to come up with the following solution.

I feel like this is the type of solution one might come up with–it is symmetric and somehow "makes sense."

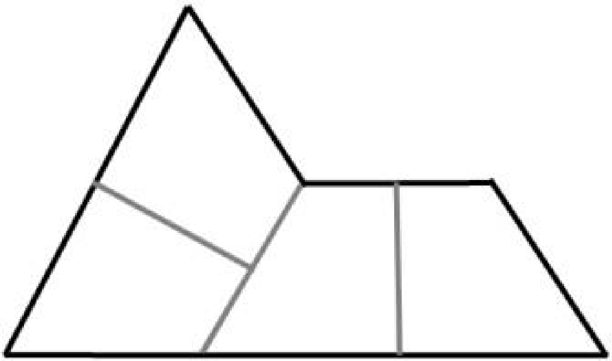

One of the students had shown a lot of creativity in his work. He came up with the above solution, but he also submitted a second answer that definitely took me by surprise.

Here's the solution that he came up with:

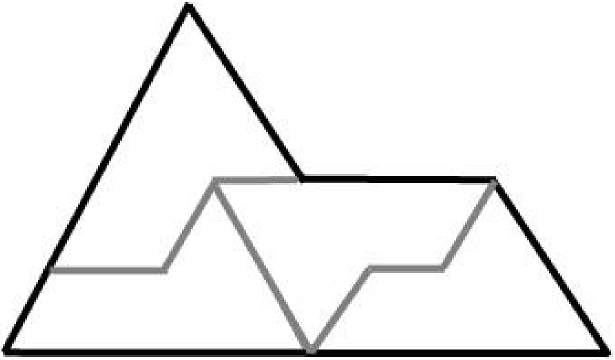

It is amazing to see how the shape can be divided into four parts using scaled down versions of itself! Well done if you came up with this answer on your own.

I had wondered what it would be like if the process was repeated: that is, if you continue to divide the subdivisions into 4 small versions of itself.

One reader of Mind Your Decisions, V Paul Smith took up the challenge and did a manual tessellation. Here is what it looks like. It's absolutely beautiful.

(Visit this page for the full-scale version of the tessellation: http://dl.dropbox.com/u/3990649/Tesselation 01.jpg)

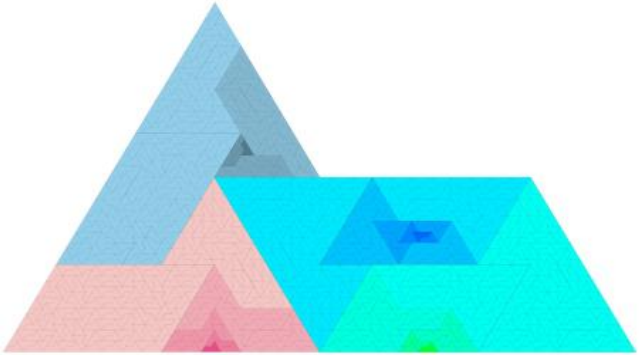

Puzzle 16: Moat crossing problem

A castle is surrounded by a rectangular moat that measures 20 feet in width: see image below.

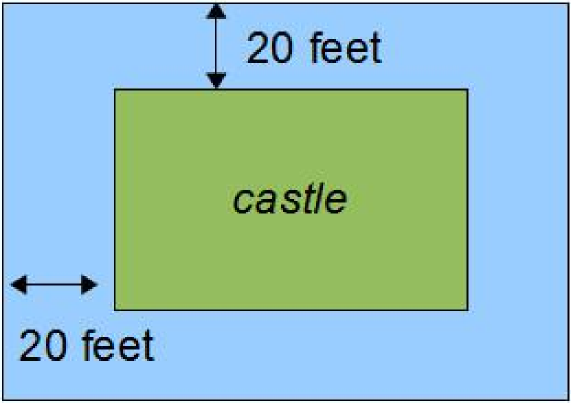

You have two planks of 19 feet each, but no way to nail them together.

How can you cross the moat if you start from the outside of the moat and want to reach the castle?

Extension: what's the largest rectangular moat you can cross from the outside with two planks of length L?

Answer to Puzzle 16: Moat crossing problem

The answer to the puzzle

You can cross the moat by arranging one plank along the corner of the moat, and putting the other on top, as follows:

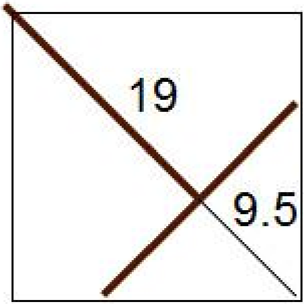

As you can check, the planks will cover a distance of 28.5 feet across the diagonal, which measures just a tad less at 28.3 feet. Thus, you have just enough plank to cross.

Using geometry, we can figure out the largest moat that can be traversed with a plank of length L.

We will solve for the longest width that corresponds to the planks in their optimal position. The answer is: $L/\sqrt{2} + (L/2)/\sqrt{2}$

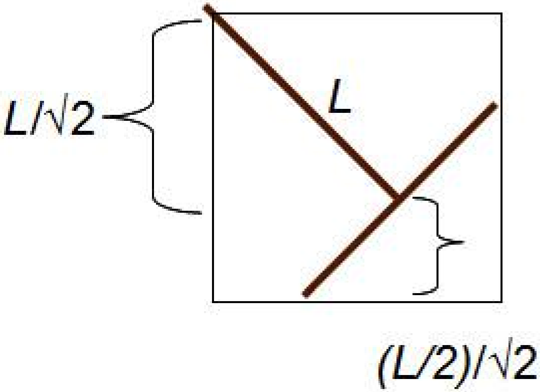
L/√2
L
(L/2)/√2

Puzzle 17: Mischievous child

At a dinner party, there are two large bowls filled with juice. One bowl holds exactly 1 gallon of apple juice and another has 1 gallon of fruit punch.

A mischievous child notices the bowls and decides to have a little fun. The child fills up a ladle of apple juice and mixes it into the bowl with fruit punch. Not content to stop here, he decides to do the reverse. He fills up a ladle of the fruit punch/apple juice mixture and returns it to the apple juice bowl.

The child would proceed further, but his mother notices what he is doing and makes him stop. The child apologizes to the hosts, who decide to shrug off the matter as little harm was done.

But an interesting question does arise about the two mixtures of juice.

In the end, the two bowls ended up with some of the other juice. The question is: which bowl has more of the other juice? That is, does the fruit punch bowl have more apple juice or does the apple juice bowl have more fruit punch?

Assume the ladle holds a volume of 1 cup and the juices were mixed thoroughly when the child transferred the juices.

Answer to Puzzle 17: Mischievous child

The hard way to solve this problem is doing the math. You can calculate that both juice bowls end up with an equal concentration of the other juice, and thus the transferred volumes must be equal.

The easier way is to think logically. Notice that both bowls begin and end up with exactly 1 gallon of liquid. This means that whatever apple juice ended up in the fruit punch bowl must have been replaced by the same volume of fruit punch that went into the apple juice bowl. Therefore, the two volumes must be equal!

The problem is known as the wine/water puzzle. If you'd like a more detailed solution, I found a nice explanation here: wine/water problem solution

Puzzle 18: Table seating order

A table seat choice can be the difference between a boring, wasted night and a fun, profitable one. I can recall two examples where seat choice made a big difference.

The first was a student-faculty dinner at Stanford where I had invited a math professor. The etiquette was to accompany a professor while getting food and walk to a table. The natural instinct, therefore, was to sit directly next to the invited professor. But this was a bad choice, as it was difficult to make eye contact and direct conversation. It also led to awkward moments where students spilled food and drinks on their professor. Lesson learned: always sit across the table!

The second came in a friendly poker game. After playing a few times, we quickly learned the importance of seating order, particularly when betting in no-limit Texas Hold'em. We have since paid careful attention to rotate seats for fairness.

The games of course led to a natural question: exactly how many different betting orders are possible? (that is, how many ways can people sit around a table, if only their relative position matters?)

Answer to Puzzle 18: Table seating order

There are a handful of ways to determine the answer. Here are a few that I like.

Method 1: Converting linear permutations into circular permutations

The case of two people is trivial: there is only one way.

How many ways can three people sit around a table? One way is to count permutations.

The easiest type of permutation to count is a "linear" list. Say the people around the table are sitting as person A, then person B, and finally person C. We can represent this order in a linear list as ABC.

Using this notation, we can count the number of possible lists. We simply note there are:

3 possible choices for the first spot (A, B, or C)
2 choices for the second
1 choice for the last spot

This means there are 3 x 2 x 1 = 6 = 3! ways to write the list. Specifically the list can be written as:

ABC
ACB
BAC
BCA
CAB
CBA

But this list is not our answer. At least some of these permutations represent the same seating order on a circular table. We can see this graphically:

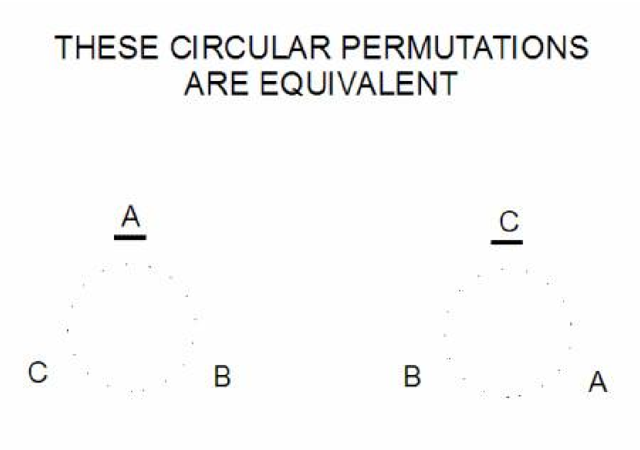

The image above shows that the list orders ABC and CAB are the same arrangement on a circular table.

So we ask: what's the relation between linear permutations and the circular ones we wish to count.

The relationship can be illustrated as follows:

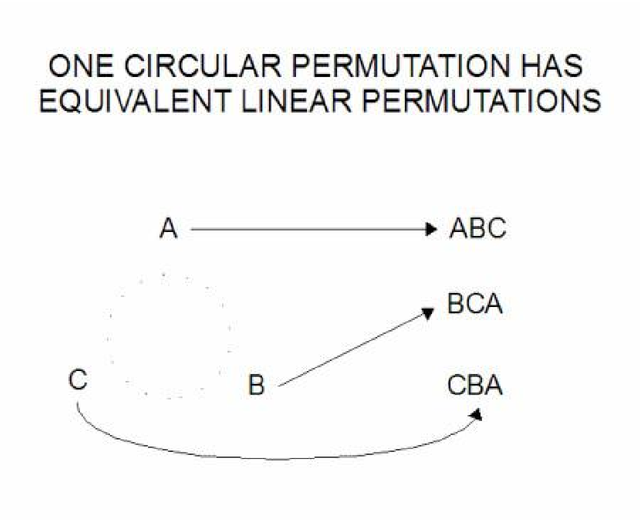

Evidently, each circular permutation for a three-person group can be written in 3 different ways. This makes sense: for each circular permutation, there are three different choices for the first letter of the linear permutation representation.

Thus, we can convert the number of circular permutations into linear ones by multiplying by 3. Or working in reverse, we can convert the number of linear permutations into circular ones by dividing by 3.

Combing all of this, we can deduce there are 3! / 3 = 2 ways to seat three people on a table. (The answers are ABC and ACB)

We can expand this logic to more people. We first count the number of linear permutations and then convert to circular ones.

For four people, the number of linear permutations can be counted. There will be 4 choices for the first spot, 3 choices for the second, 2 choices for the third, and 1 choice for the last. Therefore there will be 4 x 3 x 2 x 1 = 4! linear permutations.

We can then convert this into the number of circular permutations. As there are 4 people in the group, there will be 4 ways that each circular permutation can be written as a linear permutation–any of the four people can be written first in the list. So now to convert linear into circular we divide by 4 (again the number of people in the group).

Thus there will be a total of 4! / 4 = 6 ways to seat this group.

To generalize even further, we can see a pattern for n people. We can write the linear permutation in $n!$ ways, but we have to divide by n to convert the linear permutations into circular ones.

In the end, the formula simplifies as:

$$\text{seating orders} = \frac{n!}{n} = (n-1)(n-2)\cdots 1 = (n-1)!$$

And viola, we have our answer.

Method 2: induction

An alternate way of solving this problem is mathematical induction.

Listing out a few cases of two, three, and four suggests the general formula $(n-1)!$ Now we can prove it.

Consider a group of $n-1$ people who are about to get seating at a table in a restaurant. Let's say at the very last minute one extra person comes. How many ways can the group be seated?

By the induction hypothesis, we know there are $(n-2)!$ ways for the initial group to sit. Where can the additional person sit? For any of these $(n-2)!$ arrangements, he can obviously sit between the first and second person, or between the second and third, or so on until the last position of being between the $n - 1$ person and the first person.

This is a total of $n - 1$ spots he can sit for any of those $(n - 2)!$ arrangements. Therefore, this group of n people has this many arrangements:

$$\text{seating orders} = (n-1) \times (n-2)! = (n-1)!$$

And like mathemagic, induction proves the formula.

Method 3: seat-changing permutations

A final way I like to visualize the answer is a party-game type approach.

Consider for a moment that n people have sat at a circular table. How many ways can they *switch* seats and have at least one person sitting with different neighbors on left and right sides? This is another way of asking the number of circular permutations, so the answer to this question will answer our original question.

Some ways people switch will obviously not change the seating order. If everyone moves one seat to the right, then each person has the same neighbors and so the seating arrangement is the same.

Such rotations do not change the order of seating.

And this demonstrates a principle: motion is always relative to a reference point. To count the number of distinct seat trades, we must have a fixed reference point. Without loss of generality, we can choose any one person to be a reference point.

With one person firmly seated, every unique linear ordering of the remaining people change seats will be a unique circular permutation.

In other words, we want to know the number of linear permutations for n - 1 people. The answer is $(n - 1)!$ and we've found the answer again.

Puzzle 19: Dart game

Alice and Bob play the following game with their friend Charlie.

Charlie begins the game by secretly picking a spot on the dartboard. The spot can be anywhere on the board, but once picked it does not change.

Then Alice and Bob each get to throw one dart at the board.

At this point, Charlie reveals the position he initially picked. The winner of the game is the person whose dart is closest to the spot Charlie picked.

For example, if Charlie picked the spot marked with an “x”, and Alice and Bob shot as follows, then Alice would win the game:

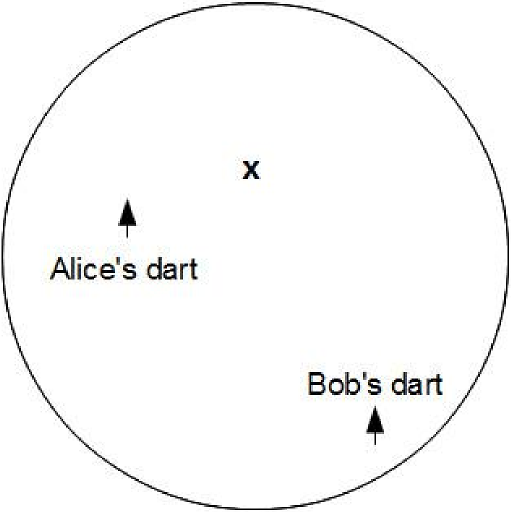

Put yourself in the shoes of Alice or Bob. What strategy is best for playing this game?

Answer to Puzzle 19: Dart game

The best strategy is fairly intuitive: Alice and Bob should each shoot for the center of the dartboard.

One way to think about this is probabilistically. Because Charlie is essentially picking a random position anywhere on the board, the best spot to pick would be the average position, which is the center of the dartboard.

Another way to think about this is geometrically. Suppose Alice hits the center of the dartboard and Bob hits somewhere else. We can ask: what is the set of all points that are closer to Alice's dart than to Bob's?

The answer can be found by remembering a fact from geometry. The set of all points that are equidistant from Alice' and Bob's darts is defined by the perpendicular bisector between the two points (the line that goes through the midpoint of the two points, is perpendicular to the line connecting the points).

The perpendicular bisector separates all the points that are closer to the different darts. All the points to one side of the perpendicular bisector must be closer to Alice's dart, and all the points on the other side are closer to Bob's.

If Alice hits the center, and Bob hits anywhere else, then the perpendicular bisector will always be some chord of the circle not going through the center. Geometrically, there will be more points closer to Alice's dart than to Bob's dart. Therefore Alice's dart "covers more ground" and she will have a higher chance of winning the game.

In the following diagram, all points to the left of the perpendicular bisector are closer to Alice's dart, and that covers more than half the board.

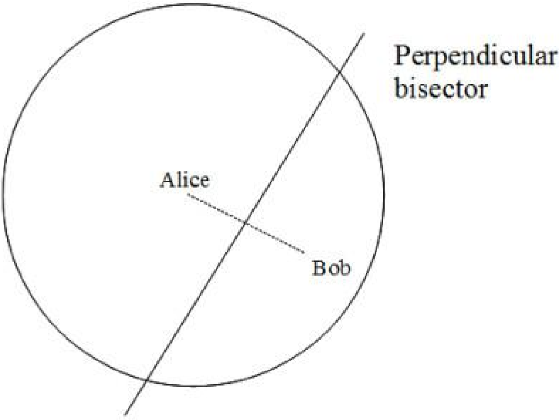

Locational games like this can prove useful in military settings or business settings when two competing parties need to position themselves closer to an unknown target (consider two hostile nations, one trying to capture and another trying to protect a terrorist hiding out in an unknown location).

This dart game is also a two-dimensional version of Hotelling's game, in which two hot dog vendors compete to locate closer to customers on a beach. In that game too it is the best strategy for each vendor to locate centrally. I explained more details in this post that shows why gas stations locate next to each other.

Puzzle 20: Train fly problem

This is another classic math puzzle.

Two trains that are 60 miles apart are headed towards each other. Each train is moving at 30 miles per hour.

A speedy fly can travel at 60 miles per hour leaves from the front of one train and heads towards the other train. When it gets to the front of the other train, it instantly turns back towards the original train. This continues until the two trains pass each other, at which point the fly stops.

The question is, how far did the fly travel?

Answer to Puzzle 20: Train fly problem

The story goes this puzzle was asked to polymath John von Neumann at a party. He quickly gave the right answer and explained he knew no trick, he just summed up the infinite series.

There is a really neat trick to solving this puzzle that does not involve infinite series.

The shortcut is to think about the problem in terms of speed and time. The distance the fly travels can then be obtained by multiplying those two quantities.

We know the fly travels at 60 miles per hour, so we have it's speed. Let's figure out the time.

The two trains are 60 miles apart, and they are traveling towards each other at 30 miles per hour each, to make for a combined speed of 60 miles per hour. Therefore, the trains will meet in 1 hour (both trains will have traveled 30 miles to the center).

Since the fly was moving for 1 hour at 60 miles per hour, that means the fly must have traveled 60 miles in all. Note this calculation ignores the actual flight path of the fly, which is precisely the trick.

Solving the problem using an infinite series is much harder: here's one derivation. I don't know anyone who could have done the infinite series using mental math–heck, it's hard enough on paper. But von Neumann's calculating abilities were so impressive that it was actually plausible.

Puzzle 21: Train station pickup

Mr. Smith, a commuter, is picked up each day at the train station at exactly 5 o'clock. One day he arrived at the train station unannounced at 4 o'clock and began to walk home. Eventually he met the chauffeur driving to the station to get him. The chauffeur drove him the rest of the way home, getting him there 20 minutes earlier than usual.

On another day, Mr. Smith arrived at the train station unexpectedly at 4:30, and again began walking home. Again he met the chauffeur and rode the rest of the way with him. How much ahead of usual were they this time?

(Assume constant speeds, and that no time is lost turning the car around and picking up Mr. Smith.)

Answer to Puzzle 21: Train station pickup

The first time I solved the problem I wrote out several equations and solved for everything algebraically. When I figured out the answer, I realized the puzzle can be solved much easier!

When Mr. Smith arrived at the train station 1 hour early, and started walking home, he was able to save 20 minutes of commute. This is due to two reasons: the driver met him closer to home (by 10 minutes), and the drive home was shorter (by 10 minutes).

So if Mr. Smith arrives 30 minutes early, or half of 1 hour, we can deduce he only traverses half the distance as before. Thus, the time savings are halved: he meets the driver closer to home by 10/2 = 5 minutes, and he drive home is shorter by 10/2 = 5 minutes.

Therefore, Mr. Smith arrives home 10 minutes ahead of schedule.

(credit: puzzle from this website)

Puzzle 22: Random size confetti

Professor X teaches a probability class. He assigns a holiday-themed project to his students.

Each student is to create a 500 rectangular-shaped confetti pieces, with length and width to be random numbers between 0 and 1 inches.

Alice goes home and gets started. She interprets the assignment as follows. Alice generates two random numbers from the uniform distribution, and then she uses the first number as the length and the second as the width of the rectangle.

Bob interprets the assignment differently. He instead generates one random number from the uniform distribution, and he uses that number for both the length and width, meaning he creates squares of confetti.

Clearly Alice and Bob will cut out different shapes of confetti. But how will the average size of the confetti compare?

That is, will the average area of the shapes that Alice and Bob cut out be the same? If not, whose confetti will have a larger average area?

Answer to Puzzle 22: Random confetti

Let X be a random variable with a uniform distribution.

Bob takes one realization of X, so the area he cuts out will be distributed as X^2, and the expected area is $E(X^2)$

Alice instead takes two realization of X. The area she cuts out will be E(X)*E(X), or $E^2(X)$.

The difference between Bob's expected area and Alice's is:

$$E(X^2) - E^2(X) = Var(X) >= 0$$

The difference between Bob's expected areas and Alice's is equal to the variance of X, which is always a non-negative number. Notice this formula holds for random variables of other distributions too, like normal distributions or discrete distributions.

In the case of the uniform distribution from 0 to 1, the variance is 1/12.

So Bob's areas will always be at least as large or larger than Alice's. So Bob may need a little bit more paper than Alice when cutting his confetti.

Puzzle 23: Hands on a clock

The long hand of a very accurate timepiece points exactly at a full minute, while the short hand is exactly two minutes away. What times of day could it be?

Answer to Puzzle 23: Hands on a clock

The trick is realizing there are limited times that the hour hand lands exactly on one of the minute markings. Since the hour hand moves from one hour number marking to the next (5 minute markings) in a span of 60 minutes, that means the hour hand is only on minute markings every 12 minutes of the hour, corresponding to the minute times 00, 12, 24, 36, and 48.

From here it is just an exercise in trial an error to figure out the right times. If the minute hand is at 00, the hour has to be near 11, 12, or 1 to solve the puzzle. But in these times the two hands are separated by either 5 or 0 markings.

For the minute hand at 12, the candidate time would have the hour nearby at the number 2. But at 2:12, the hour hand has moved one marking, and the minute hand is two markings past the number 2. The two hands are separated by just one minute marking.

We can proceed to figure out the minute hand at 24 will work. If the minute hand is at 24, the candidate hour hand would be nearby at 4. We can check 4:24 exactly works: the hour hand is 2 markings past the clock "4", and the minute hand is 4 markings past the clock "4". So does the next candidate of 7:36.

Finally, you can check that the minute hand at 48 does not work.

So the two times are 4:24 and 7:36, either am or pm.

Puzzle 24: String cutting problem

An interviewer gives you a string that measures 4 meters in length. The string is to be cut into two pieces. One piece is made into the shape of a square, and the other into a circle. (picture of this below)

Example of cutting string

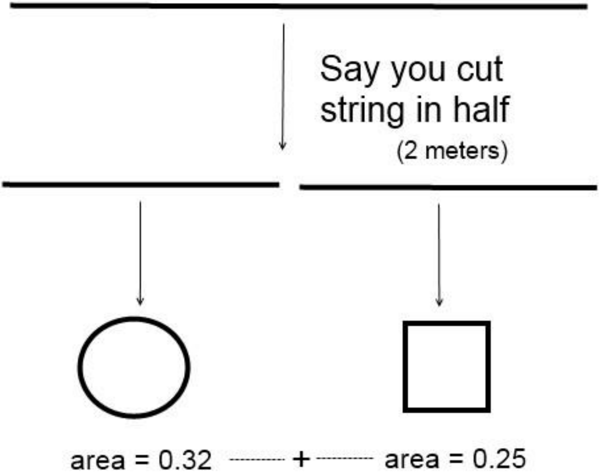

Total enclosed area = 0.57

Your job is to make the total enclosed area *as large* as possible.

The interviewer hands you a piece of paper and a pencil so you can do the math (you only get one chance to cut the string so you want to be sure your first attempt is correct).

1. How should you cut the string to maximize the area?

If you are able to figure out the answer, the interviewer has a couple follow-up questions to test your skills.

2. How should you cut the string if you want to *minimize* the enclosed area?

3. Imagine the string is cut randomly. What is the average value of the enclosed area? (When you cut the string, there is one piece to the left of the cut and another to the right. Suppose the left piece is always made into a circle and the right into a square)

Answer to Puzzle 24: String cutting problem

1. How should you cut the string to maximize the area?

This is something of a trick question. For a given length, the circle is the shape that encloses the largest area. So you want to make the whole string be the circle. (This is known as the isoperimetric inequality and it is not a trivial thing to prove!)

As you must cut it into two pieces, you should try to cut as close to one end as possible to make the rectangle small.

2. How should you cut the string if you want to minimize the enclosed area?

This can be solved using calculus. If you let one side of the string be called x for the circle and the other side be called 4 - x, then you can find a formula for the area of the two shapes as follows:

Area $= x^2/(4\pi) + (1 - x/4)^2$

The first term is the area of the circle, and the second the square.

The formula is for a parabola.

Using calculus (skipping steps here) we can find the minimum happens at $(4\pi)/(4 + \pi)$.

3. Imagine the string is cut randomly. What is the average value of the enclosed area?

As stated in step 2, the area function is described by the equation:

$f(x) = x^2/(4\pi) + (1 - x/4)^2$

We can take the average value by calculating an integral: you integrate the function from 0 to 4 (which gives the area under the curve) and then you divide by the length of the interval (4) to arrive at the average value:

> Average value = 0.25 integral (f(x), x, 0, 4)

Using calculus this is computed as 0.758

Puzzle 25: One mile South, one mile East, one mile North

This is a very classic puzzle, a fitting end.

I first read about this in the fun puzzle book How Would You Move Mount Fuji?

Years ago, Microsoft apparently used to ask this puzzle as an interview question.

Here is the problem: how many points are there on the earth where you could travel one mile south, then one mile east, then one mile north and end up in the same spot you started?

Answer to Puzzle 25: One mile South, one mile East, one mile North

This puzzle is much harder than it seems at first.

The easy but wrong answer

The place that comes to mind is the North Pole.

This is, in fact, one of the correct spots.

You can trace out the path on a globe. From the north pole, you can move your finger south one mile. From there, you will go east one mile and move along a line of latitude that is exactly one mile away from the north pole. You finally travel one mile north, and you will exactly end up in the north pole.

The route you travel will look like a triangle or a piece of pie, as seen in this rough sketch I made:

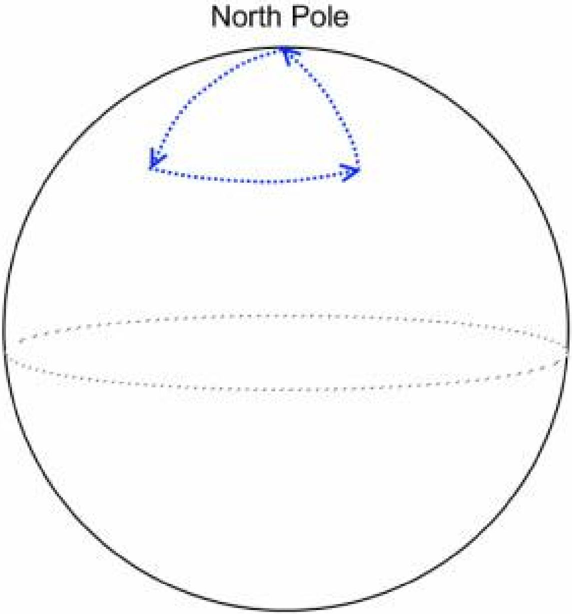

This is one correct answer. But it is not the only one.

The harder spots

The other spots on the earth all involve traveling near the South Pole.

The trick to these solutions is that you end up in the same spot after traveling one mile east.

How can that be?

One way this is possible is if you are on a line of latitude so close to the South Pole that the entire circle of latitude is exactly one mile around. We will label this circle C(1) for convenience.

With this circle in mind, it is possible to figure out a solution.

Let us begin the journey from a point exactly one mile north of C(1). Let's trace out the path of going one mile south, one mile east, and one mile north again.

To begin, we travel one mile south to point on the circle C(1). Then, we travel east along the circle C(1), and by its construction, we end up exactly where we began. Now we travel one mile north, and we reach the starting point of the journey, exactly as we wanted.

The trip will look something like this rough sketch I made:

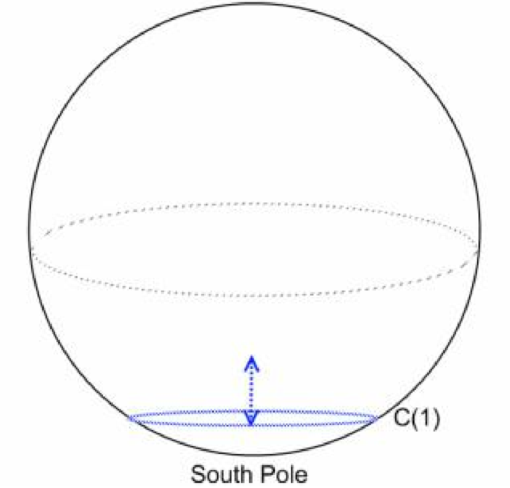

This demonstrates there is a solution involving a circle near the South Pole.

In fact, the circle C(1) is associated with a family of solutions. Any point one mile north of C(1) will be a possible solution. This means the entire circle of latitude one mile north of C(1) is a solution. This means there are an infinite number of solutions associated with the circle C(1)!

That alone seems remarkable. But what is more interesting is that there are even more solutions.

The circle C(1) was special because we traversed it exactly once, and ended where we started from, when we went one mile east.

There are other circles with the same property. Consider the circle C(1/2), a similarly defined circle of exactly 1/2 mile in circumference. Notice that traveling one mile east along this circle will also send us back to the starting point. The only difference is that we will have traversed the circle two times!

Thus we can construct solutions using the circle C(1/2). We start one mile north from C(1/2) and every point along this line of latitude is a solution. There is an infinite number of solutions associated with the circle C(1/2).

Naturally, we can extend this process to more circles. Consider the circle C(1/3), similarly defined with exactly 1/3 mile in circumference. It would be traversed three times if we travel one mile east along it, and we would end in the same place we started from. This circle too will have an infinite set of solutions–namely the line of latitude one mile north of it.

To generalize, we can construct an infinite number of such circles. We know the circles C(1), C(1/2), C(1/3), C(1/4), … C(1/n), … will be traversed exactly *n* times if we travel one mile east along them. And there are corresponding starting points on the lines of latitudes one mile north of each of these respective circles.

In summary, there are an infinite number of circles of latitudes, and each circle of latitude contains an infinite number of starting points.

The correct answer, therefore, is one point at the North Pole, plus two infinite set of points of circles near the South Pole.

Section 2: Probability problems

Life is often said to be a game of chance. The following 25 puzzles deal with probability.

Puzzle 1: Making a fair coin toss

Alice and Bob play a game as follows.

Alice spins a coin on a table and waits for it to land on one side.

If the result is heads, Alice wins $1 from Bob; if tails, Alice pays $1 to Bob.

While the game sounds fair, Bob suspects the coin may be biased to land on heads more. The problem is he cannot prove it.

Being diplomatic, Bob does not accuse Alice of trickery. Instead, Bob introduces a small change in the rules to make the game fair to both players.

What rule could Bob have come up with?

Answer to Puzzle 1: Making a fair coin toss

Bob worries the coin may be biased to land on heads more often than tails. The trick Bob comes up with is a way to turn a biased coin into having fair tosses.

The technique is referred to as the von Neumann procedure, and it works as follows:

> Step 1. Spin the coin twice.
>
> Step 2. If the two results are different, use the first spin (HT becomes “heads”, and TH becomes “tails”).
>
> Step 3. If the two results are the same (HH or TT), then discard the trial and go back to step one.

In other words, Bob has redefined the payout rule to ensure the odds are fair to both parties.

Why does the von Neumann procedure works? The procedure takes advantage that HT and TH are symmetrical outcomes and will thus have equal probability.

To see this, suppose the outcome heads occurs with probability 0.6 and tails with probability 0.4. Then we can directly calculate the probability of the pairs as:

–HT occurs (0.6)(0.4) = 0.24
–TH occurs (0.6)(0.4) = 0.24

These events are equally likely, and hence both players have an even chance of winning the game.

The von Neumann procedure takes advantage that each coin flip is an independent event, and so both mixed pairs of tosses will have equal chances.

Appendix: spinning vs tossing

Observant readers may have noticed the game is about coin spinning rather than coin tossing.

Why the distinction? It's a small bit of trivia that coin tossing is not easily biased:

> "The law of conservation of angular momentum tells us that
> once the coin is in the air, it spins at a nearly constant rate
> (slowing down very slightly due to air resistance). At any rate
> of spin, it spends half the time with heads facing up and half the

> time with heads facing down, so when it lands, the two sides
> are equally likely (with minor corrections due to the nonzero
> thickness of the edge of the coin)”
>
> via Teacher’s Corner: You Can Load a Die, But You Can’t Bias a Coin

The theory is only slightly modified in real-life. In practice, there is still a small bias towards one side of a coin.

I will refer you to this article which summarizes the results from an academic paper that points out coin flipping is almost always slightly biased.

A few of the results are:

> --"If the coin is tossed and caught, it has about a 51% chance of landing on the

same face it was launched. (If it starts out as heads, there's a 51% chance it will end as heads)"

--"If the coin is spun, rather than tossed, it can have a much-larger-than-50% chance of ending with the heavier side down. Spun coins can exhibit "huge bias" (some spun coins will fall tails-up 80% of the time)"

--A coin will land on its edge around 1 in 6000 throws, creating a flipistic singularity.

The lesson is that coin flips are better than coins being spun.

But a coin flip will still exhibit some bias, so to be fair, it may be best to use the von Neumann procedure or another choice mechanism (like a computer random number generator).

Puzzle 2: iPhone passwords

This is based on a question my friend asked me.

> Presh, real-life question for you: What is the safest way to lock my iphone?
>
> Let me explain.
>
> A friend unlocked his phone once and I grabbed it and said "so, 9,6,0, and 1, huh?" because the bulk of "tap prints" were on those numbers and, I rightly presumed, correlated to his password. He freaked out because were I a thief, I could unlock his phone pretty easily as I'd know all four numbers and that they are only used once each within the four-digit code. Not terribly safe, is it?
>
> So when setting my password, I opted to repeat a number (e.g. 1-2-3-1). That way,

someone would look at my phone and even if they could figure the three numbers I use, they would either have to guess at the fourth number (which doesn't exist) or, should they rightly figure out that I only use three independent numbers, they would have to try all possible permutations of those three different numbers within a four-digit code.

Question: is it harder to guess a password that uses only 3-digits or one that uses a 4 distinct digits?

Would it be harder to guess if one only used a passcode containing 2-digits?

Answer to Puzzle 2: iPhone passwords

We need a way of counting possible passwords. The easiest case is when someone uses 4 unique numbers for the 4-digit passcode. Each number is used exactly once in the passcode, and hence the problem reduces to counting the number of ways to rearrange 4 objects. This is solved by counting the number of permutations. For a password using 4 digits, there are exactly 4! = 4 x 3 x 2 x 1 = 24 ways to have this kind of password.

But what happens when you have a password like 1231? That is, how can you count passwords in which one or more numbers are used multiple times? You have to count the number of combinations.

The way to solve this is by using an extension of permutations known as the multinomial coefficient. The multinomial coefficient is calculated as the total number of permutations divided by terms that account for non-distinct or repeated elements. If an element appears *k* times (i.e. has a multiplicity of *k*), then the factor to divide by is *k*!

A simple example from Wikipedia's entry can illustrate. Let's say we want to figure out the number of distinct ways to rearrange the letters in the word MISSISSIPPI. There are 11 letters but some of the letters are repeated. There are 1 Ms, 4 Is, 4 Ss, and 2 Ps. The number of distinct rearrangements of the letters is the number of permutations (11!) divided by the factors for the elements accounting for their multiplicity (1! x 4! x 4! x 2!). The multinomial coefficient is thus 11 ! / (1! x 4! x 4! x 2!) = 34,650.

Am I helping myself by using three numbers in a four-digit code?

There are 4! = 24 possible ways a password can be formed from four distinct and known numbers. Will using just three numbers increase the number of possibilities?

The surprising answer is that yes, it does. It seems counter-intuitive at first so let's go through an example.

Suppose you see an iPhone where the "tap prints" are on the numbers 1, 2, and 3. How many possibilities are there for the four-digit password to unlock the phone?

There's a simple observation needed to go on. In order that three numbers are all used in a four-digit password, it must be the case that some digit is used twice. Perhaps the number 1 appears twice, or the number 2, or the number 3.

Suppose the number 1 is used twice. How many passwords are possible? We can use the multinomial coefficient to figure it out. We know the total number of permutations is 4! and we must divide by 2! to account for the number 1 being used twice. Thus, there are 4! / 2! = 24 / 2 = 12 different passwords. We can list these out:

1123
1132
1213
1312
1231
1321
2113
2131
2311
3112
3121
3211

But we are not done yet. We must similarly count for the cases in which the number 2 is used twice, or the number 3 is used twice. By symmetry it should be evident that each of those cases yields an additional 12 passwords.

To summarize, there are 12 passwords when a given number is repeated, and there are three possible numbers that could be repeated. In all, there are thus 12 x 3 = 36 passwords.

Notice there were just 24 passwords when using four distinct numbers.

This trick of using three numbers does in fact increase the set of possible passwords. While each case of three digits only gives 12 passwords, the gain to this method is that the other person doesn't know which number is repeated. And so they have to consider all possibilities which becomes 36 possible passwords.

Would it be even safer if I only mixed two independent numbers?

If three is better than four, then is two better than three?

Unfortunately it is not.

There is just not enough variety when using two numbers. The gain in ambiguity of multiplicity is simply not enough to counteract the lack of passwords.

With two distinct numbers, there are only 14 possible passwords. This is found since the two numbers either have multiplicities as (1, 3), or (2, 2) or (3, 1). We can add up the multinomial coefficients to get 4! / (1! x 3!) + 4! / (2! x 2!) + 4! / (3! x 1!) = 4 + 6 + 4 = 14.

We can also list them out:

1112

1121

1211

2111

1222

2122

2212

2221

1122

1221

2211

1212

2121

2112

In conclusion, using two numbers ends up reducing the possible number of passwords.

Additional ways to help

If that weren't enough, my friend actually brainstormed a couple of other ways to improve the password.

> Actually now I can think of all kinds of brilliant maneuvers… like using three digits but tapping a phantom fourth number once the code is entered…. so there are four “tap prints” but only three which are relevant!
>
> Or, by the same measure, you could use four independent numbers and then tap a fifth time to have 5 options for four spaces.

I think these are interesting possibilities too, but they hit me as a little less practical since you’d have to diligently tap those extra numbers to make the smudge marks.

I’ll leave it to you to figure out how many passwords those methods will yield.

Perhaps an equally valuable suggestion is to simply clean the touch-screen intermittently to erase the finger print marks and leave no clue.

Puzzle 3: Lady Tasting Tea

The problem is based on an incident at a 1920s tea party in Cambridge. The story goes one lady claimed the ability to distinguish between tea made by pouring tea to milk or by adding milk to the tea.

Everyone questioned the claim, but one person decided test it out. He created an experiment with 8 tea cups, consisting of 4 cups of each preparation.

The lady was remarkably able to identify all 8 cups, raising the issue of whether she just got lucky.

What are the odds that the lady identified all 8 cups by chance?

The problem is known as the Lady Tasting Tea, and it brought about the more modern analysis of testing random experimental data.

Answer to Puzzle 3: Lady Tasting Tea

This is a classic question of combinations. We know there are 8 items, of which 4 will be one type, and 4 the other.

Therefore there are "8 choose 4" or 70 different combinations that are possible. (See the iPhone password puzzle for an explanation of counting possibilities with repeated elements)

The probability of identifying all of the cups by chance is a mere 1/70 (around 1.4 percent). We can conclude the lady likely had a refined palette.

Puzzle 4: Decision by committee

Imagine you face a very difficult decision and there is a low probability of making the right choice. ($p < 0.5$)

What would you rather do: ask a single person to decide or instead send it to a three-person group where the majority choice wins? Assume the three-person committee people are independently voting, with each having the same chance of determining the correct decision.

Answer to Puzzle 4: Decision by committee

The situation can be modeled using probability. We can say that each person has an independent probability p of making the right choice. Since the problem is difficult, we will say $p < 0.5$. (Imagine each person is equally likely to choose among three or more possible alternatives).

What's the success of the individual versus the group?

The individual is easy: the probability of making the right decision is p.

The three-person group is a little harder. The group will find the right answer whenever two or more of the people vote for the right option. Since each person can vote "right" or "wrong," there are 8 possible ways to vote:

RRR
RRW
RWR
RWW
WRR
WRW
WWR
WWW

The first, second, third, and fifth items listed are the 4 ways the group can come to the right decision. Adding the probabilities for these events gives the chance the group will come to the correct decision.

When all three are right, RRR, that is p^3. When two are right, say RRW, that is $p^2(1-p)$ and there are three such events like this.

The probability that a majority find the right decision is the sum of these events, which is $p^3+3p^2(1\text{-p}) = 3p^2 - 2p^3$

Since $p < 0.5$, we can see this final expression is less than p. In the following chart, the dotted line shows the probability the committee comes to the right decision is less than the probability an individual finds the right decision.

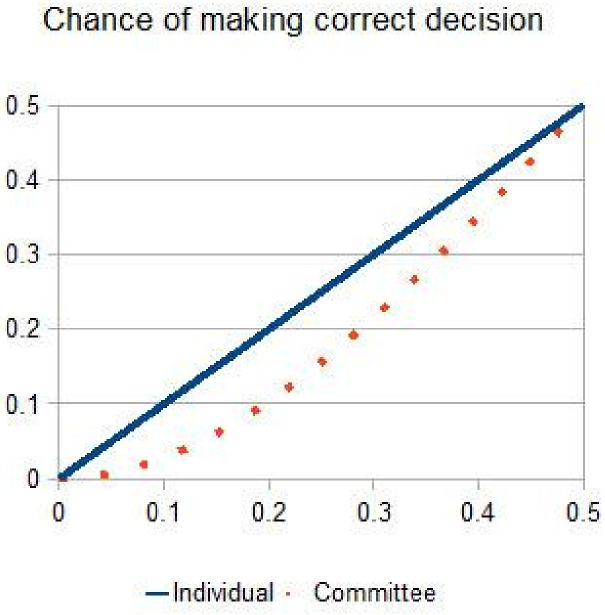

The moral: committees may not be the best for making tough choices!

That's not to say committees are useless. They will of course exist to diffuse risk and for the purpose of brainstorming (which may increase the odds of success over an individual). But this does show committees are ill-suited for the type of hard problem they are meant to address.

Puzzle 5: Sums on dice

With two dice, you can roll a 10 in two different ways: you can either roll 5 and 5, or you can roll 6 and 4. Similarly, you can roll a sum of 5 in two different ways: as the rolls and 1 and 4, or as 2 and 3.

But the two events "roll a 10" and "roll a 5" will not occur with equal frequency. Why not?

(credit: Luck, Logic, and White Lies)

Answer to Puzzle 5: Sums on dice

The trick is all about the wording of the puzzle which creates a mystery where there is none.

The sum 10 can be obtained in three ways by dice roll: namely (5,5), (4,6), and (6,4); the sum 5 in four ways: (1,4), (4,1), (2,3) and (3,2).

So the sum 10 is obtained with probability 3/36 versus the sum 5 with probability 4/36.

Pictorially:

	1	2	3	4	5	6
1				5		
2			5			
3		5				
4	5					10
5					10	
6				10		

The puzzle demonstrates that it’s always important to consider the events in probability. Sly wording, like this puzzle’s ways describing sums rather than pairs of rolls, can easily confuse.

Puzzle 6: St. Petersburg paradox

You are offered an unusual gamble.

A fair coin is tossed until the first heads appears, which ends the game. The payoff to you depends on the number of tosses. The payoff starts at 2 dollars and doubles on each successive toss.

That means you get 2 dollars if the first toss is a head, 4 dollars if the first toss is a tails and the second is a heads, 8 dollars if the first two tosses are tails and the third is a head, and so on. In other words, you get paid 2^k where k is the number of tosses for the first heads.

Toss # of first heads	Probability	Payout
1	1/2	$2
2	1/4	$4
3	1/8	$8
4	1/16	$16
5	1/32	$32
k	$(1/2)^k$	$\$2^k$

The question to you is how much should you be willing to pay to play this game? In other words, what is a fair price for this game?

Answer to Puzzle 6: St. Petersburg paradox

The typical way to answer this question is to compute the expectation (or the "average") of the payouts. This is done by multiplying the various payouts by their probability of occurrence and adding it up. To say it another way, the payouts are weighted by their likelihood.

The respective probabilities are easy to compute. The chance the first toss is a heads is 1/2, the chance the first toss is a tails and the second is a heads is 1/2 x 1/2, and the third toss being the first heads is 1/2 x 1/2 x 1/2, so the pattern is clear that the game ending on the *k* toss is $(1/2)^k$

So with probability 1/2 you win 2 dollars, with probability 1/4 you win 4 dollars, with probability 1/8 you win 8 dollars and thus the expectation is:

$$
\begin{aligned}
E &= \frac{1}{2}\cdot 2 + \frac{1}{4}\cdot 4 + \frac{1}{8}\cdot 8 + \cdots \\
&= 1 + 1 + 1 + \cdots \\
&= \infty
\end{aligned}
$$

The surprising result is the expectation is infinity. This means this game–if played exactly as described–offers an infinite payout. With an astronomical payout, a rational player should logically be willing to pay an astronomical amount to play this game, like paying a million dollars, a trillion dollars, and so on until infinity.

The fair price of infinity is paradoxical because the game does not seem like it is worth much at first. Few would be willing to pay more than 10 dollars to play this game, let alone 100 dollars or 1,000 dollars.

But expectation theory seemingly says that any amount of money is justifiable. Banks should be willing to offer loans so people could play this game; venture capital firms should offer more money than they do to start-ups; individuals should be willing to mortgage their house, take a cash advance on their credit card, and take a payday loan.

What's going on here? Why is the expectation theory fair price so different from common sense?

It turns out there are a variety of explanations.

Resolution 1: Payouts should be realistic

Imagine you are playing this game with a friend. You hit a lucky streak. The first nine tosses have been tails and you're still going. If the tenth toss is a heads, then you get

1,204 dollars as a payout. If it's a tails, you have a chance to win 2,408 dollars, and even more.

At this point your friend realizes he's made a mistake. He thought he'd cash out with your 10 dollar entry fee, but he now sees he cannot afford to risk any more.

He pleads with you to stop. He'll gladly pay you the 512 dollars you've earned–so long as you keep this whole bet a secret from his wife. What would you do in this situation?

Most of us would take the cash and show some mercy here. There is no joy in winning if it means crippling a friend financially. And this concocted scenario leads to one of the unrealistic assumptions of the St. Petersburg paradox.

In the hypothetical coin game, you're supposed to believe the other side can pay out infinitely large sums of money. It doesn't happen often, but if you get to 20 coin tosses, you fully expect to be paid 1,048,576 dollars.

This is unrealistic if you're playing with a friend or even a really, really rich friend. It might be possible with a casino, but even a casino may have a limited bankroll.

The truth is that payouts cannot be infinite. If such a game were to exist in our reality, there must be a maximum, finite payout.

This means the expectation is not an infinite sum but rather a finite sum of several terms. Depending on the size of the bankroll, the St Petersburg gamble has a finite payout.

I will spare you the details, but here are a few examples of the expectation when using a maximum payout using numbers from a Wikipedia table for illustration):

Backer	Bankroll	Expected value
Friendly game	$100	$4.28
Rich	$1,000,000	$10.95
Very Rich	$1,000,000,000	$15.93
Bill Gates (2008)	$58,000,000,000	$18.84
U.S. GDP (2007)	$13.8 trillion	$22.79
World GDP (2007)	$54.3 trillion	$23.77

(small note: these calculations are based on payouts of 1, 2, 4, etc so it's slightly different than the game I set up of 2, 4, 8, etc.)

As you can see, expectation theory now implies the fair price of the game is something like 25 dollars or less. With a more realistic model of the game, the expectation result matches common sense.

This should settle matters for anyone concerned with reality and practice, but there are people who don't accept this explanation. Such philosophers think an infinite payout is possible and so the paradox still exists.

So for these people, I will offer the following alternate resolutions that don't rely on limiting the bankroll.

Resolution 2: diminishing marginal utility

A quantity like 1,000 dollars has meaning to most people. If you were to ask a friend for such a loan, they would ask how you can pay it back, what you would use it for, and so on.

But there are times when 1,000 dollars seems to lose its value. I like to think about the show Deal or No Deal where contestants play a multi-stage lottery to win 1,000,000 dollars. At various points in the game the contestants can either keep pursuing the big prizes or they can accept smaller consolation prizes. As the prospect of a big prize increases, the contestants start to care less and less about smaller prizes like 1,000 dollars.

This is an example of the famous concept of diminishing marginal utility–the idea that at larger levels of consumption, incremental units are worth less. The concept is applicable for wealth decisions because at some point incremental earnings mean less to a person.

What this means for the St. Petersburg Paradox is that the payouts should be altered. The payouts should not be measured in dollars but rather as the utility that wealth will provide.

One way to model this is to use a logarithm function. Instead of saying the payout for the first toss is 2 dollars, we will say it is log(2) units of utility, and accordingly for the other payouts.

Using a log utility function, the St Petersburg game now has a finite payout. Here is my hand-written derivation:

$$E = \frac{1}{2}\log 2 + \frac{1}{4}\log 4 + \frac{1}{8}\log 8 + \cdots$$

$$= \frac{1}{2}\log 2 + \frac{2}{4}\log 2 + \frac{3}{8}\log 2 + \cdots$$

[Let $u = \log 2$]

$$E = \frac{1}{2}u + \frac{2}{4}u + \frac{3}{8}u + \cdots$$

$$-\quad \frac{1}{2}E = \frac{1}{4}u + \frac{2}{8}u + \cdots$$

$$\frac{1}{2}E = \frac{1}{2}u + \frac{1}{4}u + \frac{1}{8}u + \cdots$$

$$\frac{1}{2}E = u \sum_{i=1}^{\infty} \frac{1}{2^i}$$

$$= u$$

$$E = 2u$$

$$= 2\log 2 < \infty$$

This is a small payout but the actual quantity does not matter: it is just that the payout is less than infinity, showing again, that there is no real paradox here.

Puzzle 7: Odds of a comeback victory

Consider two teams A and B that are completely evenly matched. Given that a team is behind in score at halftime, what is the probability that the team will overcome the deficit and win the game?

Assume there are no ties, and the result of the first half does not affect how players perform in the second half (that is, the first and second half are taken to be independent events).

(credit: problem based on page 11, "[Probability: the language of randomness](#)," by Jeffry S. Simonoff)

Answer to Puzzle 7: Odds of a comeback victory

Because the teams are evenly matched, you might mistakenly think the answer is 50 percent. But that is the probability the team would win overall. If a team is down at halftime, the chances of winning will be less. So let us try to calculate the odds.

We have to think about how a team could have a comeback victory if it is down at halftime.

Let us first write down the possible outcomes of the game, broken down by halves. Since the two teams are evenly matched, there are four different possibilities for who is leading during each half (ignore the case of a tie):

> (first half, second half):
> AA

AB
BA
BB

Because the teams are evenly matched, these events are all equally likely so each occurs with probability 1/4 = 25 percent

In two of the cases, one team scores more points in both halfs of the game, and there is no come from behind victory: AA and BB. This means 50 percent of the games the team that lags behind at half ends up losing the game.

The other two possibilities are times when a team could have a comeback victory. In these cases, one team leads at the half, but gets outscored by the other in the second half: AB and BA. In order for a team to get a comeback victory, it must overcome the deficit from the first half. How often does that happen?

The answer can be calculated by the following logic: since the two teams are evenly matched, it is equally likely that the team will score enough points to overcome the deficit or that it will not score enough points. (For instance, the event of falling behind 6 points in one half happens with the same probability of gaining 6 points in a half). Therefore, in the event AB, it will be equally likely that B scores enough to eventually win, or that it would not score enough and it loses.

Therefore, B ends up winning in half of the cases, or 12.5 percent of the time (take 1/2 of 25 percent). The same logic applies for the event BA: there is a 12.5 percent chance that team A ends up winning.

Putting this all together, we have:

> Probability(team having comeback victory) = P(AB)*Pr(B wins) + Pr(BA)*Pr(A wins) = 12.5 + 12.5 = 25 percent

So under these assumptions, a team will have a 1 in 4 chance of making a comeback victory.

Now you may point out this is not realistic as the model does not take into account quality of teams and things like home field advantage. Nor does it take into account psychology: a recent study shows that teams with a slight deficit at halftime end up winning more often than teams with a slight edge at halftime. Here is the remarkable study based on 18,000 professional basketball games and 45,000 college games.

However, even though the assumptions are a bit off, the overall league statistics seem to mirror the probability model.

In the National Football League, a small sample of games in 2005 showed this trend:

> "Joe Gibbs is not telling his troops they have a 23 percent chance of winning. Of the 88 games observed, 68 of the teams that went in at halftime with the lead went back to the locker room at the end of the game with the lead and the win. That's right 77 percent of the time if a team had a lead at halftime, it won the game. [And thus 23 percent of the time, the team facing a deficit came back for a victory]"

I found the same pattern was shown to happen in the National Basketball League (though granted this is 20 year old data; I'd love to see whether the pattern holds true for more recent seasons):

> "Professor Hal Stern of the University of California at Irvine examined 493 National Basketball Association games from January to April 1992, and found that in 74.8% of the games, the team

> that was ahead at halftime ultimately won the game [and thus the losing team at halftime came back with probability 25.2 percent]"

This is either a pure coincidence or there is something to be said about the simple probability model. It's fascinating to me either way.

Puzzle 8: Free throw game

Alice and Bob agree to settle a dispute by shooting free throws.

The game is simple: they take turns shooting, and the first one to make a shot wins.

Alice makes a shot with probability 0.4 while Bob makes his shots with 0.6.

To compensate for the skill difference, Alice gets to shoot first.

Is this a fair game?

Extension: if Alice makes a shot with probability p and Bob with probability q, for what values of p and q would the game be fair? Solve if $q = 1 - p$

Answer to Puzzle 8: Free throw game

There are many methods to solving the probabilities of winning. The one I like is a technique of backwards induction.

The free throw game seems hard to figure out because a round could end with no one making a shot, and then the game would continue. Solving for the winning probability seems like you'd need to use an infinite series.

But that's not the case. The trick is seeing that each round is really an independent sub-game. The fact that the previous round ended without a winner does not affect the winner of the current round or any future round. This means we can safely ignore outcomes without winners.

The probability of winning depends only on the features of a single round.

This simplifies the problem to a more tractable one. So now, assume that one of the players did win in a round, and then calculate the relative winning percentages.

We can use a little trick to visualize the problem. Because Alice makes a shot with probability 0.4, and Bob with 0.6, we can imagine the two are not shooting free throws but instead rolling a 5 sided die.

Let's say that Alice makes her shot for rolling the numbers 1 and 2, and Bob makes his shot for the other three numbers 3, 4, and 5.

So Alice wins if she rolls one of her winning numbers. If she does not, then Bob gets a chance to roll and he wins the game for rolling his numbers. All other combinations of the rolls means they both miss their shots, so the round is a draw and they go again.

Here is a diagram illustrating the outcomes of a round, illustrating the events for which Alice will win:

				Bob makes his shot		
		1	2	**3**	**4**	**5**
Alice makes her shot	**1**	W	W	W	W	W
	2	W	W	W	W	W
	3	-	-	L	L	L
	4	-	-	L	L	L
	5	-	-	L	L	L

To calculate the winning percentage, we can simply count out the number of ways that Alice wins. In the grid, there are 10 squares that she wins, and only 9 that Bob wins.

Therefore, Alice wins with probability 10/19 = 53 percent and Bob only with 9/19 = 47 percent.

Although Bob is a better shooter, Alice has a slight edge in the game because she gets to shoot first.

Answer to extension: generalizing the probabilities

The numbers we used made it convenient to convert the game into rolling a 5-sided die.

But we can generalize the process.

Notice that Alice won on 0.4 percent of the squares, which is the same as her shooting percentage.

The percentage of squares for when either person won the game was 0.76, which is equal to the chances Alice makes her shot

(0.4) or she misses her shot (1 - 0.4) and Bob makes his (0.6). The sum of this is 0.4 + (1 - 0.4) * 0.6.

Thus the probability Alice wins a game is: (SP for shooting percentage)

(Alice's SP) / (Alice's SP + (1 - Alice's SP) * Bob's SP)

If we say that Alice's SP is p and Bob's is q, then this becomes:
$p / (p + (1 - p) * q)$

The game is fair if this term equals 0.5. Skipping some of the algebra, this simplifies to:
$(p - q) -pq = 0$

We can plot out all values for which this equation is true, remembering that p and q are probabilities so they are between 0 and 1. The dotted line corresponds to setting the condition $q = 1 - p$

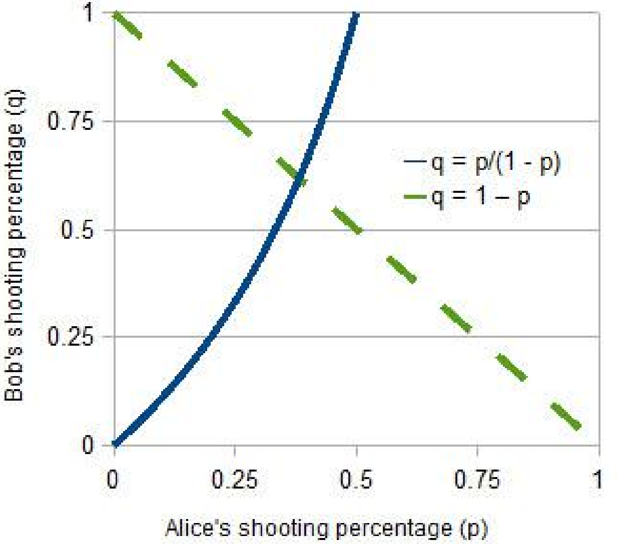

If we give the additional restriction that $q = 1 - p$, then we can uniquely solve that p is about 0.382, which is plotted above.

So Alice at 0.4 shooting percentage is just a tad higher than the fair shooting percentage of 0.382.

Puzzle 9: Video roulette

Bob loves the TV show Law & Order. Each day he picks an episode at random and watches it. Given there are 456 episodes of the show, how many days will it take Bob to watch the entire series on average?

Extension: Figure out a formula for a show that has n episodes.

Answer to Puzzle 9: Video roulette

Consider smaller cases to get an idea.

If a series has just 1 episode, it will take 1 day to watch the entire series.

What about 2 episodes? On the first day, Bob will watch one of the episodes. How long will it take him to watch the remaining episode on average?

We can solve for the number of days N as a sum of two conditional events. If he picks the episode he has not seen (with probability 0.5), then the conditional expectation is 1 day. If he instead picks the episode he has seen, then he essentially loses a day, and he is back to the starting point–so the expectation is $N + 1$.

In other words,

N = 1 * Pr(picks episode he has not seen) + (N + 1) * Pr(picks episode he has seen)

$N = 1 * 0.5 + (N + 1) * 0.5 = 0.5\,N + 1$

$N = 2$

Note that it takes Bob 2 days on average to watch the unique episode that he picks with probability 1/2.

Thus, it takes Bob an average of 3 days (1 day for the first episode, 2 days for the second) to watch a series with two episodes.

Solution

We can think about the problem in terms of rolling a die. Each day Bob picks a new episode randomly is essentially like Bob rolling a die where each face represents an episode number.

The question is: how many times on average must a 6-sided die be rolled until all sides appear at least once?

The first roll can be any of the faces. On the second roll, there are 5 remaining unique faces out of 6. Using the logic above, we can conclude it will take an average of 1 / (5/6) = 6/5 rolls until one sees a different face.

We continue the logic to calculate the number of rolls until a new face. As there are 4 remaining out of 6, this will take 6/4 rolls on average. Continuing the logic, we can conclude the total number of rolls it will take on average to reveal every face at least once is:

$1 + 6/5 + 6/4 + 6/3 + 6/2 + 6/1 = 147/10 = 14.7$

In other words, for a series with 6 episodes, it will take Bob about 15 days to watch every episode.

For a series with n episodes, the similar series is:

$$1 + \frac{n}{n-1} + \frac{n}{n-2} + \cdots + \frac{n}{1} = \sum_{i=1}^{n} \frac{n}{i} = n \sum_{i=1}^{n} \frac{1}{i}$$

For n = 456, this sum is roughly 3,056.

For very large n, the series sum is roughly n* ln(n)–though this approximation for 456 yields 2792 so it is a very rough approximation.

Puzzle 10: How often does it rain?

In Mathland, the weather is described either as sunny or rainy, nothing in between.

On a sunny day, there is an equal chance it will rain on the following day or be sunny. On a rainy day, however, there is a 70 percent chance it will rain on the following day versus a 30 percent chance it will be sunny.

How often does it rain in Mathland, on average?

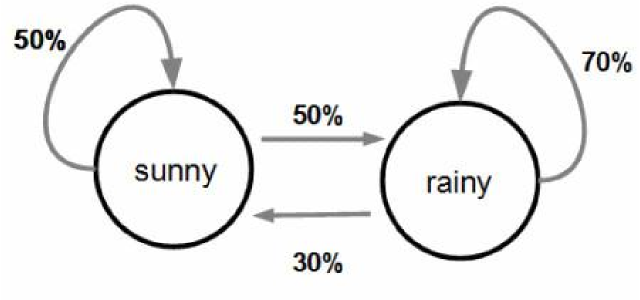
50%
70%
50%
sunny
rainy
30%

Answer to Puzzle 10: How often does it rain?

Here are a couple of ways I solved this problem.

Method 1: Let R denote the average proportion of rainy days and S of sunny days. Using the law of total probability, we know that:

R = E(rains tomorrow | sunny today) * Pr(sunny today) + E(rains tomorrow | rainy today) * Pr(rainy today)

We can know use a clever trick. On average, the probability it is sunny or rainy on a particular day is S and R, respectively. And we also know a day is either sunny or rainy, so S = (1-R). Hence we get the following:

R = E(rains tomorrow | sunny today) * (1-S) + E(rains tomorrow | rainy today) * R

We can simplify this because we know the rules of weather. The first conditional expectation is 50 percent and the second is 70 percent.

$R = 0.5(1-R) + 0.7(R)$

This can be solved to find out $R = 0.625$.

Method 2: The weather can be modeled as a regular or ergodic Markov chain. This is far beyond the scope of this puzzle book, so see section 11.3 of this pdf for a reference on long-run averages.

From that reference it is shown that for a regular transition matrix A, there is a unique row vector w called the probability vector such that wA = w. This is the probability of being in those states.

Let us set up the transition matrix:

$$A = \begin{bmatrix} 0.7 & 0.5 \\ 0.3 & 0.5 \end{bmatrix}$$

If w = (R S), and we set wA = w, we get the equations:

$$0.7\,R + 0.5\,S = R$$
$$0.3\,R + 0.5\,S = S$$

These equations can be solved to find R = 0.625, just as before.

Puzzle 11: Ping pong probability

Suppose A and B are equally strong ping pong players. Is it more likely that A will beat B in 3 out of 4 games, or in 5 out of 8 games?

(credit: problem in this math book)

Answer to Puzzle 11: Ping pong probability

It is more likely that A will beat B in 3 out of 4 games than in 5 games out of 8. It's a simple application of binomial probability distributions.

The equation for A winning exactly r games out of n is "n choose r" times 0.5^n

We can calculate the chance of A beating B in exactly 3 games out of 4 is 25% and the odds of A winning exactly 5 games out of 8 equals 7/32 or roughly 21.8%.

Why is this the case? The counter-intuitive part is the wording of the question. If we instead consider A to win at least 3 out of 4 games, or at least 5 out of 8 games we have a different situation.

In that case, we find that the probability of winning 3 or more games out of 4 increases to 31.25%. All we need to do is add the probability of winning exactly 3 games out of 4, which equals 25%, to the probability of winning all 4 games, which is (.5)^4= .0625, or 6.25%.

The probability of winning 5 or more games out of 8 is equal to (7/32)+ .1094 +.0312+.0039 = .3632, or 36.32%.

So in this problem it's important that we want to know A is winning exactly a certain number of games, rather than winning at least some number of games.

Puzzle 12: How long to heaven?

A person dies and arrives at the gates to heaven. There are three identical doors: one of them leads to heaven, another leads to a 1-day stay in limbo, and then back to the gate, and the other leads to a 2-day stay in limbo, and then back to the gate.

Every time the person is back at the gate, the three doors are reshuffled. How long, on the average, will it take the person to reach heaven?

Answer to Puzzle 12: How long to heaven?

Let N denote the average number of days it takes to get to heaven.

The trick to solve for N is to rewrite the average using symmetry of the game.

N is equal to the average number of days regardless of which door you enter first. This splits up into three cases:

Case 1: One-third of the time you go directly to heaven, and that's 0 days.

Case 2: One-third of the time you pick the door that adds 1 day. In this case, you end up in heaven in $N + 1$ days.

Case 3: The remaining one-third you pick the door that adds 2 days. In this case, you end up in heaven in $N + 2$ days.

These observations lead to the following equation and answer:

$$N = 1/3 * 0 + 1/3 * (N + 1) + 1/3 * (N + 2)$$

$$N = 1/3 * N + 1/3 + 1/3 * N + 2/3$$

$$N = 2/3 * N + 1$$

$$1/3 * N = 1$$

$$N = 3$$

Puzzle 13: Odds of a bad password

This is a problem that I was asked by a reader.

The problem is as follows: A system has 100 accounts, two of which have bad passwords (let's call these bad accounts). If someone could only test 20 accounts, what are the chances that one will net a bad account?

Extensions:

1. What is the probability of netting both bad accounts in the sample of 20? What about exactly one bad account?

2. What is the probability of netting a bad account if you have k bad accounts, there are N total accounts, and you can sample n accounts at one time?

3. Go back to the problem with 100 accounts, and 2 bad accounts. Suppose you can vary how many accounts you can sample. If you want a 50 percent chance of netting a bad account, what's the minimum sample size needed?

Answer to Puzzle 13: Odds of a bad password

The answer can be found using the hypergeometric distributrion.

The same problem can be restated as follows: if you are drawing 20 balls from an urn of 98 white balls and 2 black balls, what are the chances of drawing a black ball?

The way I calculated this is to find the chance of drawing only white balls and finding the complement event. Thus the chance is:

> 1 - (2 choose 0)(98 choose 20)/(100 choose 20) = 36 percent.

1. What is the probability of netting both bad accounts in the sample of 20? What about exactly one bad account?

Two bad accounts is:

> (2 choose 2)(98 choose 18)/(100 choose 20) = 19/495 or about 4 percent

One bad account is:

> (2 choose 1)(98 choose 19)/(100 choose 20) = 32/99 or about 32 percent

2. What is the probability of netting a bad account if you have *k* bad accounts, there are *N* total accounts, and you can sample n accounts at one time?

You find the chance of getting no bad accounts and then find the complement.

This is:

> 1 - (k choose 0)([N - k] choose n)/(N choose n)

3. Go back to the problem with 100 accounts, and 2 bad accounts. Suppose you can vary how many accounts you can sample. If you want a 50 percent chance of netting a bad account, what's the minimum sample size needed?

I used a numerical method to vary the sample size and found out the answer is 30.

It's interesting that you only need to sample about a third of the population to have a better than even chance of finding both bad accounts!

Puzzle 14: Russian roulette

Can probability theory save your life? Perhaps not in usual circumstances, but it sure would help if you found yourself playing a game of Russian roulette.

Let's play a game of Russian roulette. The rules are this: I have a gun that has six empty chambers. Now watch me as I put a *single bullet* in the gun. I close the cylinder and spin it. I point the gun to your head and, click, it turns out to be empty.

Now I'm going to pull the trigger one more time and see if you are really lucky. Which would you prefer, that I spin the cylinder first, or that I just pull the trigger?

(credit: I'm not sure of the original source of this puzzle, but the wording is similar to a problem in William Poundstone's book How Would You Move Mount Fuji?

Answer to Puzzle 14: Russian roulette

The problem can be solved by calculating the probability of survival for the choices.

First, consider the odds of survival if the cylinder is spun. The cylinder is equally likely to stop at any of the six chambers. One of the chambers contains the bullet and is unsafe. The other five chambers are empty and you would survive. Consequently, the probability of survival is 5/6, or about 83 percent.

Next, consider the odds if the cylinder is not spun. As the trigger was already pulled, there are five possible chambers remaining. Additionally, one of these chambers contains the bullet. That leaves four empty or safe chambers out of five. Thus the probability of survival is 4/5, or 80 percent.

Comparing the two options it is evident that you are slightly better off if the cylinder is spun.

Puzzle 15: Cards in the dark

You are given pack of cards has 52 cards in a completely dark room. Inside the deck there are 42 cards facing down, 10 cards facing up.

Your task is to reorganize the deck into two piles so that each pile contains an equal number of cards that face up. Remember, you are in the darkness and can't see.

How can you do it?

Answer to Puzzle 15: Cards in the Dark

Take any ten cards from the original deck. Create a new deck by flipping over each card one by one. The two decks will contain the same number of cards facing up.

Why is that?

Verifying this is a relatively simple counting exercise. Suppose, for example, the 10 cards you took consisted of three face up cards and seven face down cards. Since every face down card gets flipped in the new deck, the new deck will therefore consist of 7 face up cards. This exactly matches the original deck which has 7 remaining face up cards (since three face up cards were removed for the new deck).

The idea is this: removing a card and flipping it is a matching action. When you remove a face down card in the original deck, the number of face up cards is unaffected, which is matched by the new deck getting a face up card. When you remove a face up card, the number of face up cards is subtracted by one, which is matched by the new deck getting a face down card. By repeating the matching action ten times (the number of cards facing up in the original deck), you guarantee that both the new deck and the old deck will have the same number of face up cards.

The general proof goes something like this. Of the 10 cards you remove, suppose the number of face up cards removed is x. That leaves the original deck with $10 - x$ face up cards. Correspondingly the new deck contains those x cards with a face down orientation. Thus the remaining $10 - x$ are face up cards and the two decks match.

(You can extend the problem too. If the original deck had 15 face up cards, then you create a new deck by choosing 15 cards and flipping them over. The proof is analogous.)

This puzzle generated a lot of comments online, incidentally. My favorite comment: "The existentialist's solution: throw all the cards in the trash and make two piles of zero."

Puzzle 16: Birthday line probability

During a probability course, the professor announces a chance for the students to get extra credit.

First, the students are to form a single-file line, without knowing the rules of the game.

Then, the professor announces the rule. The person who gets the extra credit is the first person to have a matching birthday of someone in front of them in the line.

The poor first person has no chance of winning. But which person in line has the best chance of winning? What is that probability?

Assume birthdays are distributed uniformly across the year, and the students formed the line randomly because they did not know the rules.

(credit: adapted from website braingle)

Answer to Puzzle 16: Birthday line probability

The answer is the 20th person in line has the best chance of winning at 3.23%.

The puzzle can be solved analytically, and the algebra is written over at braingle.

But this is in fact a perfect problem to use a numerical method. Here is how I solved the problem.

Clearly the first person in line has 0 percent chance of winning. The second person in line wins if he matches the first person, which happens with probability 1/365. Let's denote the probability the second person wins as p(2).

What about the third person? He can only win if two things happen. One, the first two people cannot have matching birthdays. This probability is (1 – p(2)). Two, he has to match one of the two previous birthdays. Since the first two did not match, there are 2 possible birthdays the third person could have.

Putting this together, we have

$p(3) = (1 - p(2)) * 2/365$

We can generalize this formula. The probability the fourth person wins is the probability the first three people did not win times the probability he matches any of 3 birthdays. So the probability the n^{th} person wins is equal to:

$p(n) = (1 - p(2) - \ldots - p(n-1)) * (n - 1)/365$

This recurrence relation is easily programmed into a spreadsheet. Have one column that lists the position of the person n, another that has the formula (n – 1)/365, and a final column for the cumulative sum of winning probabilities for the n– 1 people ahead in line.

Here is an illustration of the probability distribution:

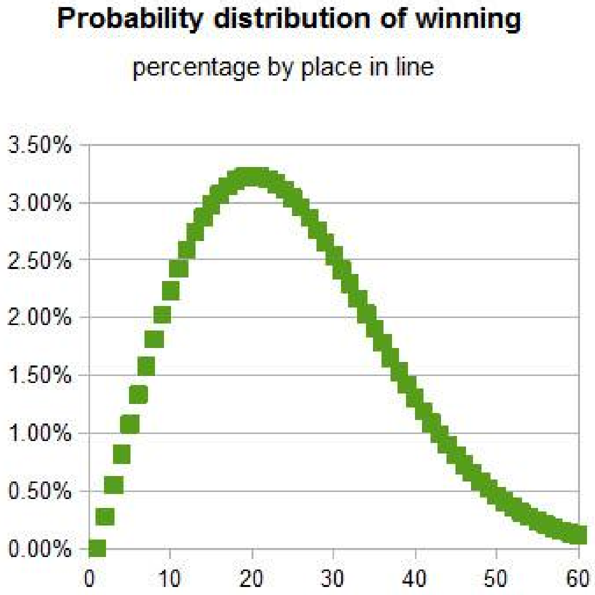

The peak probability happens at position 20, with value 3.23 percent. Nearby positions like 19 and 21 are almost the same probability, so in this case it does help if you are close to the correct answer.

Puzzle 17: Dealing to the first ace in poker

In Texas Holdem poker, sitting in the dealer position is a strategic advantage. The dealer position generally acts last in betting and is not forced to post blinds.

For a game in progress, the dealer position rotates around the table after each hand. But at the start of the game, the dealer position is simply assigned to one player.

So who gets to be dealer initially?

In poker tournaments, the dealer position is chosen by a random process so everything is fair.

In home games, people do not always use random number generators. One of the common methods is dealing to the first ace. It works like this: the host deals a card to each player, face up, and continues to deal until someone receives an ace. This player gets to start the game as dealer.

The question is: does dealing to the first ace give everyone an equal chance to be dealer? Is this a fair system?

Answer to Puzzle 17: Dealing to the first ace in poker

Dealing to the first ace is not a fair system. The distribution of the first ace appearing on the k^{th} card is not completely random.

To solve the problem, it is helpful to solve a related question: what is the probability that the first time an ace is dealt from the deck is the 1^{st}, or 2^{nd}, or 3^{rd}, or the k^{th} card?

Once this distribution is known, it will then be possible to calculate the odds a person will get the dealer by summing up the possible "winning" positions. But more on that later. For now, let's calculate the probability distribution of the first ace being dealt in position *k*.

To begin, note that a standard deck has 52 total cards of which 4 are aces.

What are the odds an ace will be the 1st; card dealt? The probability is readily calculated as the number of aces divided by the total cards which is 4/52.

So far easy enough. Continuing, what are the odds an ace will first be dealt as the 2nd card from the deck? This happens only if the following two events occur:

(i) the first card dealt was not an ace (48/52)
AND
(ii) the second card dealt is an ace (4/51)

I have written the probabilities at the end of each condition. The probability for (i) is the number of non-ace cards divided by the number of total cards, or 48/52. The probability for (ii) is similarly calculated but just slightly more complicated. The numerator is the number of aces which is obviously 4. The denominator is the number of cards still left in the deck. As one card was dealt for event (i), there are 51 cards remaining. And hence the probability for (ii) becomes 4/51.

Therefore, the probability for the first ace being dealt as the 2^{nd} card from the deck is the product of these two events, which is (48 x 4) / (52 x 51).

We can continue the exercise to calculate the first ace appearing on the 3^{rd} card. This only happens when three events occur:

(i) the first card dealt was not an ace (48/52) AND
(ii) the second card dealt was not an ace (47/51) AND
(iii) the third card dealt is an ace (4/50)

The probabilities for each event are calculated in the same fashion as above: the only tricky part is remembering to decrement the numerators and denominators to account for the cards already dealt out.

Putting these together, the probability for the 3^{rd} card being the first ace is (48 x 47 x 4) / (52 x 51 x 50).

By now it is evident the probability calculation has a pattern. We can thus generalize the logic to calculate the first ace appearing on the k^{th} card.

The specifics for this to happen are the following events:

(i) the first card dealt was not an ace (48/52) AND
(ii) the second card dealt was not an ace (47/51) AND
...
(k) the k^{th} card dealt is an ace [4/(52-k + 1)]

This calculation is straight-forward and again the only tricky part is the diminishing numerators and denominators.

Multiplying these event probabilities together yields the chance as [48 x 47 x ... (48 -*k* + 2) x 4] / [52 x 51 x 50 x (52 – *k* + 1)], for 1 > *k* > 49

There is a restriction on *k* because the process can theoretically continue until there are just 4 cards left in the deck, all of which are all aces. And then the next card must be an ace.

I went ahead and calculated the probability for each *k* and I thought a graph would be instructive. Here is what the distribution looks like:

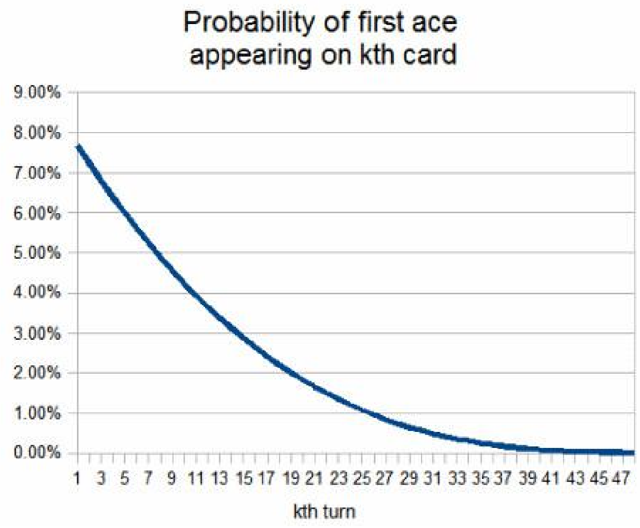

The distribution is very gently falling because it is less and less likely it will take so many turns for the first ace to appear.

Solving the original problem

Now that we have the complete distribution, we can solve for the probability a particular player is assigned the dealer.

To see how this works, consider a poker game with just two players. Let's say the first person dealt a card face up is "player 1″ and the other person is "player 2."

When will player 1 be dealer? Player 1 is dealer if the first ace is dealt to him and not player 2. Which cards are potentially dealt to player 1? Player 1 gets the first card, then one card goes to player 2, but then he gets the third card, and so on. In other words, player 1 is the dealer precisely if the first ace appears in the odd-numbers positions 1, 3, 5, ..., 47. And correspondingly, player 2 is the dealer if the first ace appears in any of the even-numbered positions 2, 4, 6, ..., 48.

Using a spreadsheet it is easy enough to sum up those entries to find the probabilities. It turns out that player 1 gets to be dealer almost 52 percent of the time versus 48 percent for player 2. This might seem like a small edge, but realize this is worse of a bias than most casino games! Player 1 has a great advantage in this system.

The entire distribution

Similar calculations can be performed if the game starts with a different number of players. For illustration, I extended the calculations for games of 3 players up to 9 players (a full ring game).

The probability is again calculated based on the distribution of the first ace. In a 3-handed game, for example, the first person dealt is the dealer if the first ace appears on the turns 1, 4, 7, etc.; the second on turns 2, 5, 8, etc.; and the third on turns 3, 6, 9, etc. (This calculation was automated as I used a handy spreadsheet array formula to sum up the probabilities based on the turn modulo).

Here are the results.

Players	Fair odds	Actual odds 1st player dealer	Edge
2	50%	52%	2%
3	33%	36%	3%
4	25%	28%	3%
5	20%	23%	3%
6	17%	20%	3%
7	14%	18%	4%
8	13%	16%	4%
9	11%	15%	4%

The first table is about the probability the first player receiving a card gets the ace. Notice there is a definite edge over the fair odds of anywhere from 2 to 4 percent.

The advantage becomes exceeding large in games with 7, 8, and 9 players where the first person receiving a card has almost double the chance of getting to be dealer initially compared to the last player.

Puzzle 18: Dice brain teaser

You and I play a game where we take turns rolling a die. I win if I roll a 4. You win if you roll a 5.

If I go first, what's the probability that I win?

Here are some clarifying notes about the game:

–If I don't get a 4, and you don't get a 5, we keep rolling until one of us does get a winning number.

–The order of play matters. If I roll a 4, I win and the game ends. You roll only if I fail to get a 4.

–Someone will ultimately win the game (there is no draw). This means the probability I win is the same as the probability that you lose.

Answer to Puzzle 18: Dice brain teaser

This dice problem is mentally tricky because many rounds end without a winner. It would seem necessary to keep track of an infinite series to arrive at an answer.

But that's not the case. The trick is seeing that each round is really an independent sub-game. The fact that the previous round ended without a winner does not affect the winner of the current round or any future round. This means we can safely ignore outcomes without winners.

Method 1: conditional probability

The probability of winning depends only on the features of a single round.

This simplifies the problem to a more tractable one. So now, assume that one of the players did win in a round, and then calculate the relative winning percentages.

In other words, calculate the probability the first player wins given the round definitely produced a winner.

To do that, we look at the distribution of outcomes. In any given round, the first player can roll six outcomes, as can the second player. How many of those thirty-six outcomes produce a winner, and how many are from the first player?

This diagram illustrates the answer:

	1	2	3	4	5	6
1					L	
2					L	
3					L	
4	W	W	W	W	W	W
5					L	
6					L	

There are exactly 11 outcomes where somebody wins, of which 6 belong to the first player. Therefore, the first player wins with a 6/11 chance, or about 54.5 percent of the time.

The first-mover advantage is caused by the fact the first player wins even if both were to roll winning numbers.

But there are a couple of other ways to think about the problem too. The next method is especially interesting.

Method 2: Symmetric Thinking

This method is about using a mental trick. I particularly find it satisfying, though I can honestly say I would not have come up with this on my own.

Here is the solution from a comment by Mja at Reasonable Deviations:

Let p denote the probability that "I" (the first player) win. Since all games ultimately produce a winner, the second player wins with the complementary probability 1 - p.

Let’s figure out the chance that I win. On my roll, I have a 1/6 chance of winning the game. What happens in the 5/6 of cases when I don’t win?

If I don’t win, the second player gets a chance to roll. Now, it’s the other person that gets to roll first and I have to wait.

This means if I do not win on my first roll, the game is the same but I take on the role of the player that rolls second.

Hence, if I do not win on my first roll, my winning chances become $1 – p$.

Algebraically, this can be written as:

p = Pr(win 1st roll) + Pr(not win 1st roll) Pr(win | not win on 1st roll)

p = Pr(roll 4) + Pr(not roll 4)Pr(second person rolling wins)

$$p = 1/6 + 5/6\,(1 - p)$$

$$p = 1 - 5p/6$$

$$11p / 6 = 1$$

$p = 6/11$, or about a 54.5 percent chance

I can't think of immediate applications besides puzzles, but the symmetry is beautiful. This is one of those solutions seeking problems.

Method 4: Infinite Series

This is the conventional solution method for math classes. It works, but I certainly find the other methods to be more interesting.

To start, we draw an infinite the game tree illustrating the outcomes for each round of the game:

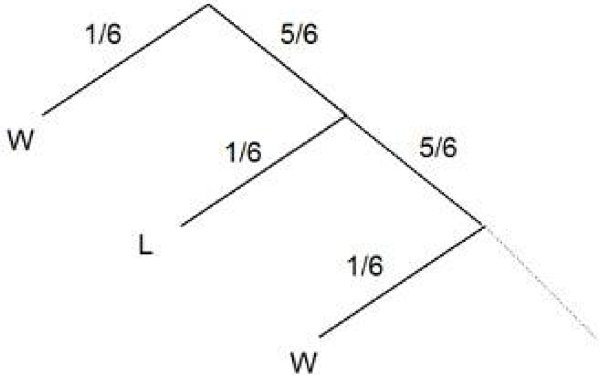

The first player's winning percentage is the sum of all branches that lead to a win. These are all the odd-numbered branches in this diagram.

The first branch is reached with probability 1/6, the third branch is reached with probability 1/6 times 5/6 squared, and each subsequent odd-branch has an extra factor of 5/6 squared.

The task is solving the following infinite series:

$$p = \sum_{i=0}^{\infty} \left(\frac{1}{6}\right) \left(\frac{5}{6}\right)^{2i}$$

Using the formula for geometric series, the solution is:

$$p = \frac{(1/6)}{1 - (5/6)^2} = \frac{6}{11}$$

So we again arrive at 6/11, or about 54.5 percent, just as before.

Puzzle 19: Secret Santa math

Suppose N people put their names into a hat, then they all draw a name. The draw is successful if no one draws his or her own name. How likely is that?

The puzzle was perhaps inspired by Secret Santa, a gift exchange in which everyone draws a name to give a gift to. The assignment is legal if no one draws his own name (gifting to oneself is not much fun). The puzzle is alternately stated as: how likely is it that a Secret Santa draw is a permissible assignment?

Answer to Puzzle 19: Secret Santa math

Let's try to figure out a pattern from analyzing a few small cases.

A bit of notation can help. We can arbitrarily label the people with numbers 1, 2, ..., n. Further, we can think about a draw as a permutation of these numbers.

I'll use the following shorthand (which is standard permutation notation). If there are three people, for example, then the notation 312 means "the first person drew name 3, the second person drew name 1, and the third person drew name 2."

Let us consider the case of $n = 3$. The possible number of draws is the number of permutations of three items, or $3! = 6$. How many of these draws are permissible–that is no one chooses his own name? We can directly list these out:

231
312

You can verify these are the only two derangements. Thus, we can conclude for $n = 3$ that the probability is

Probability(3) = 2 / 3! = 2 / 6 = 1/3 = .33333....

From this case we have figured out a solution method. We need to find the number of permissible solutions and divide it by the number of total draws (which will be equal to n!)

So what happens with four people? How different is the probability then?

The total number of draws is 4! = 24. The number of permissible draws can be figured out by direct counting, and we find there are 9 of them:

2143, 2341, 2413,
3142, 3412, 3421,
4123, 4312, 4321

So this time the calculation is:

Probability(4) = 9 / 4! = 9 / 24 = 1/3 = .375

It's interesting the probability did not change by very much from the case of n = 3.

It would be unwise to proceed in higher and higher cases by direct counting. We already have a formula for the total number of cases in the denominator (n!). What we need is a formula for the number of permissible cases in the numerator.

The general solution

How many ways can a set of objects be rearranged such that no object remains in its initial position?

There is a special name for this kind of permutation. It is known as a derangement. Also, because of its relation to permutations, there is a special notation for derangements. A derangement of n objects is abbreviated as $!n$ with the exclamation point appearing before the symbol.

Counting the number of derangements can be done in a couple of ways, beyond this book's scope, and their detailed proofs are here.

I will mention that the method I prefer uses the inclusion-exclusion principle. The idea is to count the total number of permutations (n!) and then subtract out any permutation that fixes one or more points. The tricky part is to make sure you don't double count, which is where the inclusion-exclusion formula comes in. The formula specifies how to add and subtract various subsets (like fixing one point, two points, three points, etc) so that the resulting figure does not double count.

Using the inclusion-exclusion formula, the formula for the number of derangements is:

> !n = Total permutations – permutations fixing 1 point + permutations fixing 2 points .... +/- permutations fixing *n* points

which can be simplified as

$$n!\left(1-\frac{1}{1!}+\frac{1}{2!}-\frac{1}{3!}+\cdots\pm\frac{1}{n!}\right)=n!\sum_{k=0}^{n}\frac{(-1)^k}{k!}$$

The interesting part is the summation term. This is familiar as the partial sum for the

Taylor series expansion of 1/*e* – quite an interesting development!

Thus, the number of derangements is well-approximated by !n =n! / *e* (the exact formula is !n = floor(n! / *e* + 1/2) since you need to round up for even numbers and round down for odd numbers).

So we can now solve the puzzle as *n* approaches infinity. The probability is:

Probability(n) = !*n* /*n*! ~ (*n*! / *e*) / (*n*!) = 1/*e* = 0.3679

It's again remarkable that *e* appears in a probability problem, and the answer is roughly 37 percent which is reasonably high.

Puzzle 20: Coin flipping game

Let's play a coin flipping game. You get to flip a coin, and I'll pay you depending on the result.

Here are the rules:

–you first flip a coin and we record the outcome (H or T)

–you keep flipping until the first outcome is repeated, ending the game

–you get paid $1 for each time you flipped the opposite outcome

–for instance, if you flip *H* first, and the sequence of tosses ends up as *HTTTH*, you will get paid $3 for the three *T*'s that appeared. If you flip HH, by contrast, then you will get $0

–analogous payout rules apply if you flip *T* first: if you flip *TT* you get \$0, but if you flip *THHHHT* you will get \$4

(an equivalent way of saying this is if you make a total of n tosses, you get paid $n - 2$ dollars because you don't get paid for the first or final flips)

I am going to offer you a chance to play this game for 75 cents. But there is one catch: I admit I may have biased the coin, so heads appears with probability p which may or may not be 1/2. (though it is not a two sided coin, because at $p = 0$ or $p = 1$, you obviously lose the game every time)

Should you be willing to play this game? Why or why not?

(credit: the puzzle is a problem from one of my college math books, Apostol Calculus Volume II)

Answer to Puzzle 20: Coin flipping game

It turns out I was being charitable, and you should definitely play the game. As derived below, the expected value of the game is \$1. The remarkable part is this is true regardless of the value of the chance of getting a heads p!

Let's calculate the expected value of the game.

To begin, let's write out a table of possible outcomes to the game, split up by whether the first toss is an H or a T

Toss	Probability	Payout	Expected Value
HH	p^2	0	0
HTH	$p^2 (1-p)$	1	$p^2 (1-p)$
HTTH	$p^2 (1-p)^2$	2	$2 p^2 (1-p)^2$
HTTTH	$p^2 (1-p)^3$	3	$3 p^2 (1-p)^3$
...			
HT....TH	$p^2 (1-p)^n$	n	$n p^2 (1-p)^n$
Toss	**Probability**	**Payout**	**Expected Value**
TT	(1 – p)^2	0	0
THT	$(1-p)^2 p$	1	$(1-p)^2 p$
THHT	$(1-p)^2 p^2$	2	$2 (1-p)^2 p^2$
THHHT	$(1-p)^2 p^3$	3	$3 (1-p)^2 p^3$
...			
TH...HT	$(1-p)^2 p^n$	n	$n (1-p)^2 p^n$

The expected value will be the sum of all of these individual outcomes.

We can conveniently break the game down into two contingencies, each of which is an infinite sum:

$$E(\text{game}) = E(\text{game}|H \text{ first toss}) + E(\text{game}|T \text{ first toss})$$
$$= \sum_{n=0}^{\infty} np^2(1-p)^n + \sum_{n=0}^{\infty} n(1-p)^2p^n$$

This is going to be a tricky infinite series to evaluate. We will need to use a neat trick. Note that for $0 < x < 1$

$$\sum_{n=0}^{\infty} nx^n = x\frac{d}{dx}\sum_{n=0}^{\infty} x^n = x\frac{d}{dx}\left(\frac{1}{1-x}\right) = \frac{x}{(1-x)^2}$$

We now substitute using $x = 1 - p$ and $x = p$ (from the two infinite series above) to find the value of the game.

$$E(game) = \frac{p^2(1-p)}{p^2} + \frac{p(1-p)^2}{(1-p)^2}$$
$$= (1-p) + p$$
$$= 1$$

The game has an expected value of 1, which is what we intended to prove.

Puzzle 21: Flip until heads

In mathland, one day the king asks all his subjects to perform an experiment.

He wants them to find a regular coin and record the result of its flips, under the following condition.

> Each person is to flip the coin until the result is a heads.
>
> So one person might record a result of *H* if he flips heads right away, but another person might record the result *TTH* if a heads came on the third toss.
>
> All of the million subjects are to send the record of their tosses to the king for analysis.

A royal counter will tally up the number of heads and tails from all the records. What do you expect the proportion of heads to tails to be?

Answer to Puzzle 21: Flip until heads

It is tempting to think there will be more tails than heads, as each person is flipping until a heads is seen. But remember that half of the people flip heads right away, which will change the outcome.

An expected value calculation will show the proportion of heads to tails is even at 50/50 for each.

Here is the calculation:

> 1/2 get a heads and stop: 0 tails
> 1/4 get a heads, then a tails: n/4 tails
> 1/8 get a heads, then a tails: 2n/8 tails
> 1/16 get a heads, then a tails: 3n/16 tails
> 1/32 get a heads, then a tails: 4n/32 tails
> ...
> Total: n heads and the number of tails is

$$\frac{n}{4} + \frac{2n}{4} + \frac{3n}{16} + ... = n$$

Therefore, we expect the same number of heads and tails in the population, so the proportion is 50/50.

(solution from here: boys and girls problem)

This question was once used as a Google interview puzzle, phrased in terms of families having a child until they had a boy and asking for the proportion of boys and girls in the population. I preferred to avoid the genetics of gender determination and so I used a coin flip.

The reason the question seems counter-intuitive is because the proportion of tails (or girls) in ONE family is equal to log(2) or about 31 percent. In fact, I did an experiment of my own to confirm this.

I ran 2,000 trials in a spreadsheet of having a child until the first boy, and I did two calculations. The first calculation is equal to the ratio of the total number of girls to the total number of children across all experiments. This is roughly 50 percent we calculated above.

For the second calculation I did the following. For each trial, I divided the number of girls by the total number of children. I then took the average across all the trials of this ratio. The calculation is to the average of the proportion of girls for each trial. This comes in at around 31 percent (it is precisely log(2)).

Result of experiment (n=2,000 trials)			
Total flips	Number of girls	Percentage of girls	Average (girls/# children)
4,015	2,015	50.19%	30.63%
		1 – log(2)	30.10%

As you can see, the proportion of girls expected in a specific family is a biased estimator of the proportion of girls in the total population.

The issue is that the average of the proportion for each family is not equal to the average proportion across all families.

It’s very important to know what is being asked for in probability questions!

Puzzle 22: Broken sticks puzzle

A warehouse contains thousands of sticks, each 1 meter long. One day a bored worker breaks each of the sticks in two, with each of the breaks happening at a random position along each stick. (random here means "uniform distribution")

There are three questions:

> (1) What is the average length of the shorter pieces?
>
> (2) What is the average length of the longer pieces?
>
> (3) What is the average ratio of the length of the shorter piece to the longer piece?

I will give a hint that questions (1) and (2) are easier to solve. It is much harder to solve (3) as it requires calculus.

Answer to Puzzle 22: Broken sticks puzzle

The first two questions can either be solved by considering symmetry, or they can be solved using calculus.

One might think the third question follows as a simple division from questions 1 and 2. This is not true! The average ratio is NOT the ratio of the averages! I'll explain why below.

Answer to (1)

By definition, the smaller piece will be less than half the length (0.5 meters).

The smaller sticks, therefore, will range in length from almost 0 m up to a maximum of 0.5 meters, with each length equally possible.

Thus, the average length will be about 0.25 meters, or about a quarter of the stick.

A more rigorous way of solving this, though less intuitive, is to set up an expectation and solve.

Suppose the stick is broken at point x, meaning the two pieces will be of length x and $1 - x$.

We can denote the shorter piece by the formula $\min(x, 1 - x)$.

Now we can solve for the average value by setting up an integral that ranges from 0 to 1. You will find this equals 0.25.

Answer to (2)

If the average of the smaller piece is 0.25, then it would only make sense the average of the larger piece is 0.75.

Answer to (3)

This is the most interesting piece of the puzzle.

If the smaller pieces average 0.25 meters, and the larger pieces average 0.75 meters, then wouldn't the ratio of the lengths be the division? That is, shouldn't the answer be 1/3 = 0.25 / 0.75 ?

The surprising result is no! The average ratio is not equal to the ratio of the averages.

This can be demonstrated by direct calculation.

If the stick is broken at point x, then the ratio of the shorter to the longer piece will depend on the value of x. When x is between 0 and 0.5, then the ratio is $x / (1 - x)$. When x is between 0.5 and 1, the ratio will be the reciprocal $(1 - x) / x$.

When these two pieces are integrated, here is the result:

$$\int_0^{0.5} \left(\frac{x}{1-x}\right) + \int_{0.5}^{1} \left(\frac{1-x}{x}\right) = 2\log(2)$$

The average ratio is not 1/3, but it is rather a bit higher at 0.386 (or more exactly 2 log(2) – 1, where we are using the natural log).

This itself is a rather surprising result: Euler's constant *e* comes out of nowhere!

It is seemingly paradoxical that the average ratio (shorter / longer) is not the ratio of the average of shorter to longer pieces.

The answer lies in the distribution of the ratio. Notice the chart for the ratio bows toward the center, or in other words, the peak at 0.5 seems a little "fat" in the following graph:

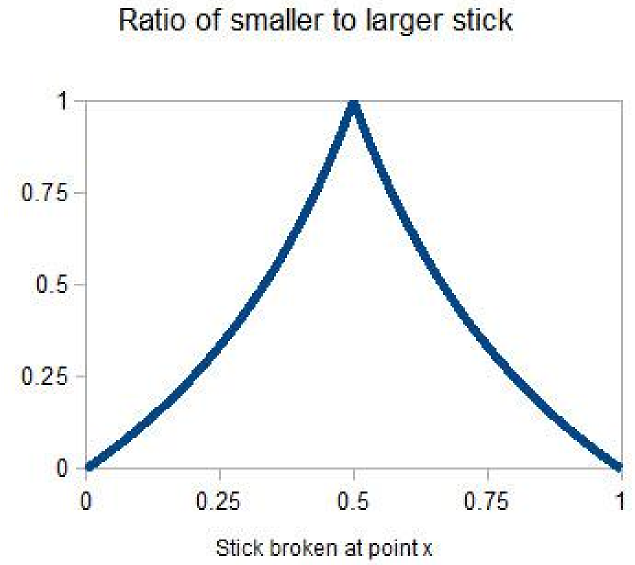

The ratio of the shorter to longer piece is slightly skewed toward the value of 0.5, and that is why the average is slightly higher at 0.386 instead of 0.333.

Puzzle 23: Finding true love

Here is a statistical model of dating.

In this statistics game, you search for your true love with sequential dates. Your only goal is to find the best person willing to date you–any thing less is a failure.

Here are some ground rules:

1. You only date one person at a time.

2. A relationship either ends with you "rejecting" or "selecting" the other person.

3. If you "reject" someone, the person is gone forever. Sorry, old flames cannot be rekindled.

4. You plan on dating some fixed number of people (N) during your lifetime.

5. As you date people, you can only tell relative rank and not true rank. This means you can tell the second person was better than the first person, but you cannot judge whether the second person is your true love. After all, there are people you have not dated yet.

How does the game play out?

You can start thinking about the solution by wondering what your strategies are. Ultimately, you have to weigh two opposing factors.

–If you pick someone too early, you are making a decision without checking out your options. Sure, you might get lucky, but it's a big risk.

–If you wait too long, you leave yourself with only a few candidates to pick from. Again, this is a risky strategy.

The game boils down to selecting an optimal stopping time between playing the field and holding out too long. What does the math say?

Answer to Puzzle 23: Finding true love

The basic advice: Reject a certain number of people, no matter how good they are, and then pick the next person better than all the previous ones.

The idea is to lock yourself in to search and then grab a good catch when it comes along. The natural question is how many people should you reject? It turns out to be proportional to how many people you want to date, so let's investigate this issue.

To make this concrete, let's look at an example for someone that wants to date three people.

Example with Three Potential Relationships

A naive approach is to select the first relationship. What are the odds the first person is the best?

It is equally likely for the first person to be the best, the second best, or the worst. This means by pure luck you have a 1/3 chance of finding true love if you always pick the first person. You also have a 1/3 chance if you always pick the last person, or always pick the second.

Can you do better than pure luck?

Yes, you can.

Consider the following strategy: get to know–but always reject–the first person. Then, select the next person judged to be better than the first person.

How often does this strategy find the best overall person? It turns out it wins 50 percent of the time!

For the specifics, there are 6 possible dating orders, and the strategy wins in three cases.

(The notation 3 1 2 means you dated the worst person first, then the best, and then the second best. I marked the person that the strategy would pick in bold and indicated a win if the strategy picked the best candidate overall.)

1 2 **3** Lose

1 3 **2** Lose

2 **1** 3 Win

2 3 **1** Win

3 **1** 2 Win

3 **2** 1 Lose

You increase your odds by learning information from the first person. Notice that in two of the cases that you win you do not actually date all three people.

As you can see, it is important to date people to learn information, but you do not want to get stuck with fewer options.

So do your odds increase if you date more people? Like 5, or 10, or 100? Does the strategy change?

The answer is both interesting and surprising.

The Best Strategy for the General Case

From the example, you can infer the best strategy is to reject some number of people

(k) and then select the next person judged better than the first k people.

When you go through the math, the odds do not change as you date more people. Although you might think meeting more people helps you, there is also a lot of noise since it is actually harder to determine which one is the best overall. So here is the conclusion.

The advice: Reject the first 37 percent of the people you want to date and then pick the next person better anyone before. Surprisingly, you'll end up with your true love 37 percent of the time.

The advice is unchanged whether you plan to date 5, 10, 50, 100, or even 1,000 people. Here is a table displaying specific numbers:

Number of people you want to date (N)	Number of people you should reject (k)
4	1
5	2
10	3
25	9
50	18
100	37

Now I was simplifying matters just a bit because "rejecting 37 percent" is an approximation. There is some math that goes into the exact answer.

To be precise, the exact answer is to find first value of *k* such that

$$\frac{1}{k+1}+\frac{1}{k+2}+\cdots+\frac{1}{N-1}-1<0$$

The problem is also known as the "Secretary problem."

The full proof is fascinating, though somewhat technical. I encourage avid math readers to check it out:

How to Find a Spouse A Problem in Discrete Mathematics With an Assist From Calculus

Puzzle 24: Shoestring problem

This is a question one of my blog readers got in an interview. It's a very hard probability puzzle to figure out on the spot.

You have a box with 30 shoe laces (or strings) in it. You can only see the ends of the strings sticking out, so you see 60 string ends total. Now you start tying them together until all ends are tied to another.

How many ways can you tie the shoe laces together?

What is the expected number of loops?

For instance there could be at least 1 big loop consisting of all the 30 strings but at most 30 individual loops when each end is tied to the end of the same string.

Answer to Puzzle 24: Shoestring problem

With math problems and interview brain teasers, there are often problem solving techniques that can help you get to the right answer.

My first thought was that working out the answer for 30 shoelaces would be hard. I would instead tackle smaller cases like considering 2 or 3 shoelaces and seeing if there is a pattern.

How many ways can you tie 2 shoelaces together?

I drew a figure like this and I counted the number of ways.

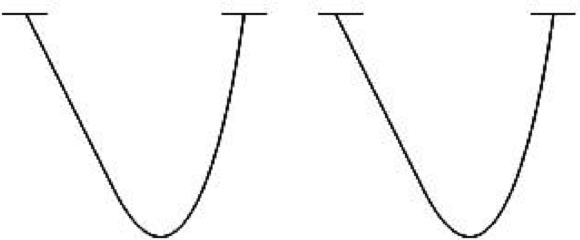

I noticed the leftmost shoelace end could connect to at most 3 spots: it could tie to the end of its own shoelace, or it could tie to one of the other two ends on the other shoelace.

After that, there would be just two ends remaining, making just 1 way to connect the loops.

This means there are exactly 3 x 1 = 3 ways to connect the ends when there are 2 shoelaces.

What about 3 shoelaces?

I tackled this case by drawing out the following diagram.

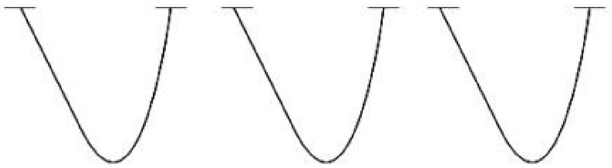

For the leftmost end, there are 5 different shoelace ends that it could connect to: either it could connect to the other end of its own shoelace, or it could connect to the four ends of the other shoelaces.

Once that end is tied, we have to consider how many ways the remaining 4 shoe lace ends could be tied. But in fact we have solved this problem already! This is exactly the number of ways 2 shoelaces could be tied together, which we found was 3 x 1 = 3.

Thus, the total number of ways for 3 shoelaces is 5 x 3 = 15.

Part 1: How many ways can you tie the shoe laces together?

We can deduce a general pattern from these cases. If there are n shoelaces, then there are $2n$ shoelace ends.

The first shoelace end can be connect to any of the other $(2n - 1)$ ends.

Tying one end to another removes 2 ends. Thus, the next shoelace end can be connected to any of the remaining $(2n - 3)$ shoelace ends, the one after that to the remaining (2n – 5) ends, and so on.

The general formula for n shoelaces is they can be tied together in:

> Ways to tie n shoelaces = 1 x 3 x 5 … x $(2n - 1)$

In other words, we have a product of the odd numbers up to $(2n - 1)$. For 30 shoelaces this will be a very big number:

> 29,215,606,371,473,169,285,018,060,091,

Part 2: How many expected loops will you get?

I spent a long time trying to figure this part out. One of the comments from this puzzle showed an amazingly elegant solution.

What about this: You have 60 ends and you start tying laces, two at a time. Each time you tie a knot, you either make a loop or you don't, and you take 2 ends out of the pool.

For each iteration, you take one end (e1) and then tie it to another end (e2). That other end is either the other end of e1 or it isn't, and you just extend the length of e1. The chance e2 is tied to the other end of e1 is 1/(n-1) where n is the number of ends in the current pool.

The expected number of loops made after the first iteration is then 1/59. For the next iteration, it's 1/57, and so on until you get to 1/1.

So the total number of expected loops is sum(1/59+1/57+...+1/1)=2.682.

Or in general, the expected number of loops is (1/1 + 1/3 + 1/5 + ... 1/(2n-1)). Quite an elegant solution!

Puzzle 25: Christmas trinkets

Assume you are running a business that sells a seasonal Christmas trinket. You can buy the trinket at $3 and sell it for $4. You can only buy the trinket once a year and cannot replenish till next year.

From experience, you have some idea about how much product will sell. Every year, the demand for the trinket from your shop will be of an equal probability between 0 and 100 (that is, there is a 1/101 chance that 0 units will sell, a 1/101 chance that 1 unit will sell, ..., and a 1/101 chance that 100 units will sell).

You have a choice to buy between 0 and 100 units of the product. After the holiday season is over, no one wants the trinkets, and you'll have to discard any unused products at your loss.

How many Christmas trinkets should you buy?

Clarification note for the probability: if you buy too few trinkets, then you simply sell out. Let's say you buy 10 trinkets, but that year the demand happened to be for 100 trinkets. In that case, you sell out of your 10 trinkets, and you missed out on the chance to profit on high demand.

(credit: this puzzle came from a question asked on Math Reddit)

Answer to Puzzle 25: Christmas trinkets

I will break the answer down into a some manageable steps.

Step 1: figuring out the probability distribution

The key to the problem is figuring out the probability distribution if you buy n units.

If you buy all 100 units, then you safely know that you have 1/101 probability of selling each unit. But what if you buy fewer units, like say 50 or 30 units? You have to derive the probability distribution from the theoretical demand.

For all units less than n, the probability that you sell that many units is simply 1/101. But for your last unit, you have to include the instances when people demand more than *n* units. To do that, you want to add in the probabilities like follows:

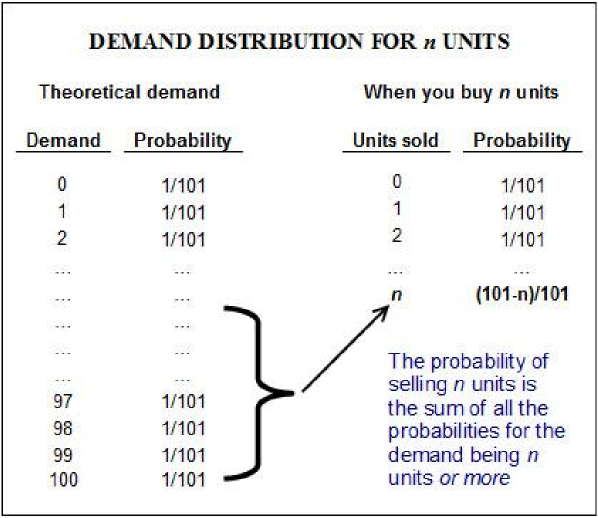

I will put the distribution in text as well for reference.

If you buy n units, then the probability you will sell units is given by:

–1/101 chance sell 0 units
–1/101 chance sell 1 unit
–1/101 chance sell 2 units
...
–(101-n)/101 chance sell all n units

The reason the last probability is higher is this: if the demand for units is higher than n, then you only get to sell n units. So you have to lump the probability of selling n or more units into one term.

Step 2: writing the expected profit

The expected profit will be given by the expected revenue (number sold times $4) subtracted by the cost (this part is easy: you spent $3 * n for the units, whether they sell or not).

So the expected profit is given by:

Profit = (selling price)(expected sales) – (cost)(units bought)
Profit = 4(expected sales) – 3n
Profit = 4[1/101 (0 + 1 + 2 + ... + n -1) + (101-n)n/101]- 3n
...
(lots of algebra)
...
Profit = 1/202 (198 n – 4n^2)

Now that you have the expected profit, the rest of the problem should be straightforward.

Step 3: maximizing profits

The amount you want to buy is the number of units that maximizes profits.

To maximize profits, we take the derivative of the profit equation and set it equal to zero.

Then we verify the amount we solved for is a maximum.

So we get:

derivative of profit = $198/202 - 8n / 202 = 0$

$n = 24.74$

Since we cannot buy fractional amounts, we need to check whether 24 or 25 is the right answer. You can find that 25 gives the maximum of $12.13 of expected profit.

In the end of the day, you ultimately want to hedge your bet and not buy too much of the supply. You buy a decent amount so you can meet demand, but if you buy too much you'll end up taking a hit on the loss.

The extension of the problem

When I solved the problem, I was curious if it meant anything that the optimal answer was buying 25 percent of the available supply.

I noticed that 25 percent was related to the margin: you make \$1 profit on a \$4 product, so that's a 25 percent margin.

This turns out to be exactly the case. Here's the general case.

Let's suppose you can sell a product for P, you buy it for C, and the available supply is S. Additionally, the demand for the product is given by:

–1/(S+1) chance sell 0 units
–1/(S+1) chance sell 1 unit
–1/(S+1) chance sell 2 units
…
–(S-n+1)/(S+1) chance sell S units

We can proceed as above to find out the expected profit of buying n units with these conditions.

I will spare you the algebra and just cut to the answer. The optimal number of units to buy is:

$$\text{optimal number} = S(1 - C/P) - 0.5 + (1 - C/P)$$

Now we have the answer, let's interpret it.

The first thing we can do is ignore the additive term $0.5 + (1 - C/P)$. Both of these are fractions, so the term will be between 0 and 1. Ultimately this will only affect the optimal answer by 1 unit, so for the sake of estimating, let's ignore this term.

So what we end up with is this:

$$\text{optimal number estimate} = S\,(1 - C/P)$$

The answer can be interpreted as follows: you should buy a percentage of supply equal to the term $(1 - C/P)$.

And what is that term $(1 - C/P)$? This is precisely the margin of the product: it's the amount of profit you make as a percentage of the price of the product.

In other words, the percentage amount of supply you should buy is equal to the margin of the product. That's quite a big simplification considering all the optimization math you see above.

Another implication of the model is this: you'll rarely want to buy all of the available supply, unless your margin is off the wall. Like if you could buy something at $1 and sell it for $100, then you're at a point where it could make sense to buy all the supply.

I love it when math works out so nicely.

Section 3: Strategy and game theory problems

Can you outthink your opponent?

The following 20 puzzles deal with strategy and game theory.

Puzzle 1: Bar coaster game

Here is how the game works:

> –Someone goes first and places a coaster anywhere on the table
>
> –The other person goes by placing a coaster anywhere else that's open on the table
>
> –The game continues with each player moving in turn to place a coaster on the table
>
> –The winner of the game is the person who puts down the last coaster, i.e., there is no more open space on the table

To make it interesting, you can play with a rule that the loser has to buy the next round.

It's a simple game, so what's the best way to play? Is it better to go first or second? Is there a winning strategy?

I don't think it matters if the table is round or rectangular, nor does it matter if the coaster is round or a square.

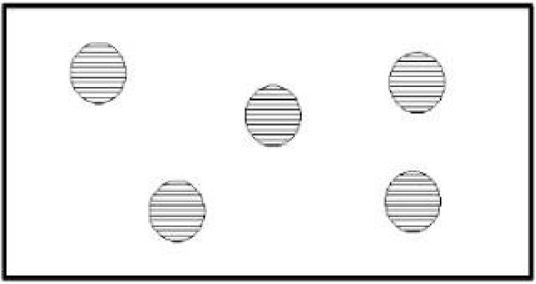

Answer to Puzzle 1: Bar coaster game

There is a winning strategy for the first player in a 2-person contest. Here is what to do:

Your first move is to place a coaster in the center of the table. Now, wherever your opponent places the coaster, you place yours symmetrically on the other side of the table. If they place a coaster in the southwest corner, you place yours in the analogous spot in the northeast corner. (If you imagine the center of the table as the origin, this is mathematically a reflection about the origin).

This strategy means you can always match your opponent's move. The game ends when your opponent runs out of open sports, which equivalently means you have placed the last coaster.

Puzzle 2: Bob is trapped

A villian has captured Bob and Alice. He could kill the dynamic duo, but he decides to have some fun while they are hostage. So he tells them they will play a game, with their lives on the line. Here is the game: Bob and Alice will be held captive in two separate facilities, under constant surveillance. Every day, each will flip a coin. Each person must then guess the result of the other person's coin (Bob has to guess Alice's toss and vice versa).

As long as one of them guesses correctly, they will both get to live for another day. But if ever both should guess wrong, then the villian will end things once and for all.

The villian smiles and then instructs the guards to proceed. Just as Bob and Alice are being taken away, Bob whispers something to Alice.

How long can Bob and Alice survive this game, on average? What must they do?

(credit: This is a puzzle adapted from Max Schireson's blog.)

Answer to Puzzle 2: Bob is trapped

Bob, super-genius that he is, devised a strategy that could allow them to survive indefinitely. The trick is the two will not be trying to guess correctly individually, but they will work as a team in their guesses.

Bob told Alice the following: every day, Alice will guess the same outcome as his flip, and Bob will guess the opposite outcome as his flip. Since the two flips will either show the same face, or the opposite face, at least one of them must be right!

To see this explicitly, here are the possible outcomes:

(Bob’s flip, Alice’s flip)

(H, H) –> same outcomes, Alice guesses correctly
(T, T) –> same outcomes, Alice guesses correctly
(T, H) –> opposite outcomes, Bob guesses correctly
(H, T) –> opposite outcome, Bob guesses correctly

Bob and Alice will keep on winning, which will no doubt provide them with enough time to devise an escape plan.

Puzzle 3: Winning at chess

Alice is a great chess player, and she occasionally taunts Bob, who barely knows the rules.

One day Bob got fed up and challenged Alice to a contest. Bob challenged Alice to play two games simultaneously, and he declared he would either win one of the games, or he would draw both of the games–in no case would he lose both games.

Bob only asked they follow a couple of ground rules. First, Bob would play black on one board and white on the other. Second, to avoid one game progressing faster than the other, they would alternate playing moves between the two boards. Bob said Alice could have the first move too.

Alice was sure she could win, and she got things going by playing her white move first.

In the end, Bob was able to draw in both games in spite of his poor skill level. How was Bob able to match wits with Alice?

Answer to Puzzle 3: Winning at chess

Alice fell right into Bob's trap, as she excitedly made the first move. On board 1, Alice opened with her move for white. So on board 2, Bob copied that move for his turn as white. Then on board 2, Alice made her reply in black. And accordingly, on board 1, Bob copied that move for his turn as black.

You can see what Bob's strategy was: he just kept copying Alice's moves for the rest of the games. Alice quickly realized that Bob was copying her moves, and that she was essentially playing against herself. If she won on one board, then Bob would surely win on the other. There was no way Bob could lose in both games. So Alice gave up and quickly moved both games into drawing positions.

Note: I suppose this strategy of copying moves could also be used for other board games of pure skill with sequential moves like connect 4, Go, checkers, etc.

One of the comments from a reader named Paul indicated this is sort of attack can happen in the real world:

> In computer security this is a well-known situation and has even been given a (gender biased) name: man-in-the-middle attack. It can be very difficult to protect against. In this chess situation, Alice would have to follow a piece of advice Presh likes to give: change the game!

Puzzle 4: Math dodgeball

Let’s analyze a math game called dodgeball that’s a sort of twist on tic-tac-toe.

Here is how the game works. It’s a two player game with the following set-up.

Player 1 gets a 6×6 grid of squares as follows:

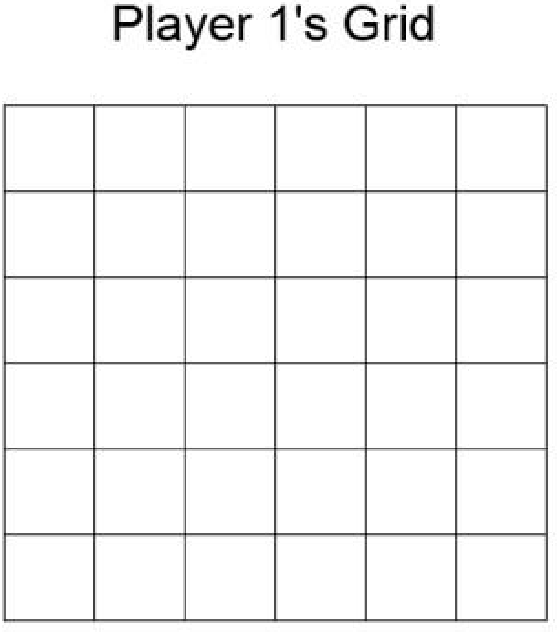

Player 2 gets a 6×1 grid of squares:

Player 2's Grid

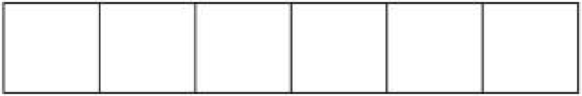

Here are the rules:

1. Player 1 begins the game by filling out the entire first ROW of his 6×6 grid, marking each square with either an X or an O.

2. Player 2 then goes by marking the just first SQUARE in his 6×1 grid, with either an X or an O.

3. On each subsequent turn, player 1 fills out the entire next row of his 6×6 grid with any combination of X's and O's. In turn, player 2 marks the next square of his 6×1 grid.

4. The game ends on the sixth round when both players have filled out their grids.

At this point, notice player 2′s grid has six squares filled with X’s and O’s. Player 1 has six such rows in his grid.

The winner is decided as follows: if player 2′s grid exactly matches one of the six rows in player 1′s grid, then player 1 wins. Otherwise, player 2 wins the game.

If you were given the choice, would you rather be player 1 or player 2? What is your strategy to win the game? Is the strategy foolproof meaning it will guarantee a victory?

Answer to Puzzle 4: Math dodgeball

It turns out player 2 can always win the game because he goes second and has an advantage.

This is not a hard game, but I will explain how it is interesting mathematically.

How can player 2 guarantee he is making a sequence that is not the same as any of player 1′s rows?

Player 2 does the following: on turn n, he looks at what player 1 writes in for the n^{th} square of the current row. Then player 2 marks exactly the opposite. For example, if player 1 begins the game by writing X in the first square, then player 2 should write O for the first square, and vice versa.

Here is an example for the 6×6 grid. After player 1 writes out a row, player 2 looks at the appropriate square and marks in the opposite. Here is what the two grids look like when the game is complete.

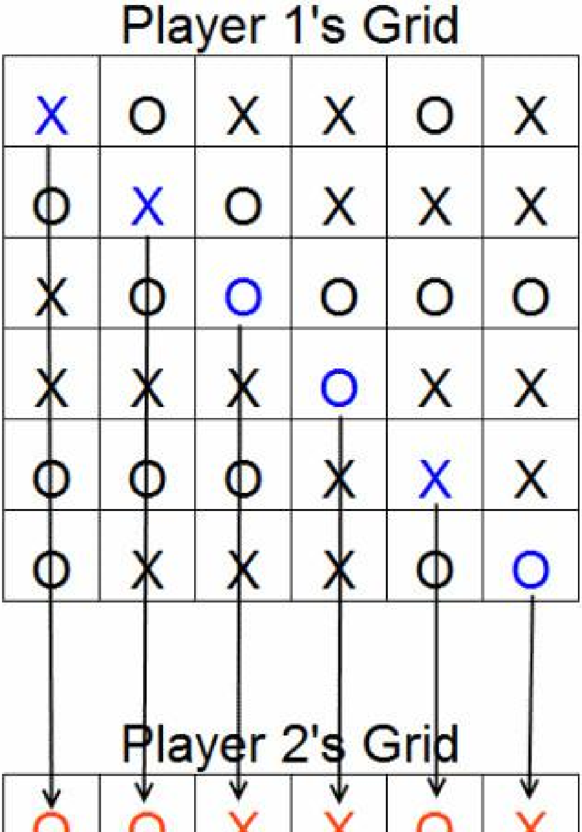
Player 1's Grid
Player 2's Grid

We can readily see that player 2′s row does not match any of the rows in player 1′s grid.

The reason is that player 2′s sequence differs from row n on the n^{th} spot, and hence the sequence must be different from any of the rows that player 1 created.

The same argument can be used to show player 2 can win for a row of any size. It even works for an infinite size grid! So even when player 1 writes an infinite number of sequences, player 2 can still make a unique sequence.

Puzzle 5: Determinant game

Alan and Barbara play a game in which they alternately write real numbers into an initially blank 1000 x 1000 matrix.

Alan plays first by writing a real number in any spot. Then Barbara writes a number in any spot, and they move in turn. The game ends when all the entries in the matrix are full.

Alan wins if the determinant of the final matrix is nonzero; Barbara wins if it is zero. Who has the winning strategy?

(credit: This problem appeared in the 2008 Putnam exam)

Answer to Puzzle 5: Determinant game

The answer is that Barbara will be able to win the game. Here is one way to see this, courtesy of the Putnam exam solutions.

The trick is to realize that the determinant of a matrix will be zero if any two rows are identically the same. Barbara can always force this to happen. How is that?

One way is for Barbara to make the first and second rows identical. If Alan writes a number in the first row, then Barbara writes the same number directly below it in the second row. If Alan writes a number in the second row, then Barbara rights the same number directly above it in the first row.

If Alan writes in any other row, then Barbara also writes a number anywhere else but rows 1 and 2.

The final matrix will have rows 1 and 2 be the same, and so Barbara will win the game.

It turns out Alan is facing a very biased game--best if he realizes this and chooses not to play.

Puzzle 6: Average salary

Three friends want to know their average salary for negotiating purposes. How can they do it without disclosing their own salaries to each other?

Answer to Puzzle 6: Average salary

The friends can calculate the average through a clever encoding process. The idea is that each person encodes their salary by adding a random number to it. These encoded salaries can be added together and then the random numbers can be subtracted. The resulting figure is the sum of the three salaries from which the average can be obtained.

Since additions and subtractions are easy to decode, however, the tricky part is implementing a solution where no person obtains knowledge of the other two party's random numbers, for that would reveal enough information to obtain individual salaries. To do that, one can sequence the additions and subtractions carefully.

Another method is more direct and works like a secret ballot, mentioned as a comment by Rajesh on my blog.

Let a pencil be a substitute for $10,000 (or some reasonable amount of money). Each person translates their salary into number of pencils, so someone making $90,000 would need 9 pencils. Then let each person drop their respective number of pencils into a sealed box. Once the three people are done, they can count the total number of pencils and find the average.

Puzzle 7: Pirate game

Three pirates (A, B, and C) arrive from a lucrative voyage with 100 pieces of gold. They will split up the money according to an ancient code dependent on their leadership rules. The pirates are organized with a strict leadership structure–pirate A is stronger than pirate B who is stronger than pirate C.

The voting process is a series of proposals with a lethal twist. Here are the rules:

1. The strongest pirate offers a split of the gold. An example would be: “0 to me, 10 to B, and 90 to C.”

2. All of the pirates, including the proposer, vote on whether to accept the split. The proposer holds the casting vote in the case of a tie.

3. If the pirates agree to the split, it happens.

4. Otherwise, the pirate who proposed the plan gets thrown overboard from the ship and perishes.

5. The next strongest pirate takes over and then offers a split of the money. The process is repeated until a proposal is accepted.

Pirates care first and foremost about living, then about getting gold. How does the game play out?

Answer to Puzzle 7: Pirate game

At first glance it appears that the strongest pirate will have to give most of the loot. But a closer analysis demonstrates the opposite result–the leader holds quite a bit of power.

The game can be solved by thinking ahead and reasoning backwards. All pirates will do this because they are a very smart bunch, a trait necessary for surviving on the high seas.

Looking ahead, let's consider what would happen if pirate A is thrown overboard. What will happen between pirates B and C? It turns out that pirate B turns into a dictator. Pirate B can vote "yes"? to any offer that he proposes, and even if pirate C declines, the situation is a tie and pirate B holds the casting vote. In this situation, pirate C has no voting power at all. Pirate B will take full advantage of his power and give himself all 100 pieces in the split, leaving pirate C with nothing.

But will pirate A ever get thrown overboard? Pirate A will clearly vote on his own proposal, so his entire goal reduces to buying a single vote to gain the majority.

Which pirate is easiest to buy off? Pirate C is a likely candidate because he ends up with nothing if pirate A dies. This means pirate C has a vested interest in keeping pirate A alive. If pirate A gives him any reasonable offer–in theoretical sense, even a single gold coin–pirate C would accept the plan.

And that's what will happen. Pirate A will offer 1 gold coin to pirate C, nothing to pirate B, and take 99 coins for himself. The plan will be accepted by pirates A and C, and it will pass. Amazingly, pirate A ends up with tremendous power despite having two opponents. Luckily, the opponents dislike each other and one can be bought off.

The game illustrates the spoils can go to the strongest pirate or the one that gets to act first, if the remaining members have conflicting interests. The leader has the means to buy off weak members.

Puzzle 8: Race to 1 million

Alice and Bob start with the number 1. Alice multiplies 1 by any whole number from 2 to 9. Bob then multiplies the result by any whole number from 2 to 9, and the game continues with each person moving in turn.

The winner is the first person to reach 1 million or more.

Who will win this game? What is the strategy?

Answer to Puzzle 8: Race to 1 million

My first attempt to solve this game demonstrated a common mistake in solving this type of puzzle. My initial attempt was to simulate the game and look for a pattern in how Bob might be able to force a certain product. This is not wrong necessarily, but it is a lot harder to see the pattern.

I then realized the game should be solved in reverse using backwards induction. The thought process is like this.

Let's imagine that we win the game, and that we are the player that sends the total beyond 1 million. The question is this: if we were able to bring the total above 1 million, what possible move could have gotten us there? That is, what number did we use to multiply the previous total by, and what previous totals would have allowed us to win?

Clearly one possibility is 500,000. If we were presented with that number, we could multiply it by 2 and get to 1 million. In fact, if we were presented with any total 500,000 or higher, then we could win by multiplying by 2. This shows that if we receive a number 500,000 or higher, then we can win the game. We can thus say that any number 500,000 or higher is a *winnable number*.

We can then ask, what other numbers are winnable numbers? In the game, we are allowed to pick any whole number from 2 to 9. Since the highest number we can pick is 9, we will use this to find the lowest number that will let us win. We can calculate that 111,112 times 9 is just over 1 million.

Therefore, any number from 111,112 to 999,999 is a winnable number.

Repeat the logic to solve the game

Now comes the interesting part. We know that if we begin a turn with any of those winnable numbers, WE win the game. The question is: what numbers came before that? What numbers, for which our opponent begins a turn, would force them to bring the total to a winnable number for us?

One example readily comes to mind: 111,111. If our opponent started with this number, the only options are to multiply it by a number from 2 to 9. The lowest total that will result is 222,222 and the highest total is 999,999. Regardless of what the opponent does, the resulting number is a winnable number for us. We can say therefore that 111,111 is a *losable number*.

What other numbers are losable numbers? We are looking for a range of numbers that will force someone to produce a result in a winnable number range. As each player has to multiply by a minimum of 2, we can find the lower range. We can calculate that 55,556 is half of 111,112.

Therefore, any number from 55,556 to 111,111 is a losable number.

Generalizing the process

We can continue reasoning. If we know the numbers 55,556 to 111,111 are losing numbers, what number range would allow a player to bring an opponent into the losable number range? These numbers would therefore also be winnable numbers.

As we reasoned above, to calculate winnable numbers we divide the lower bound by 9, and to calculate losable numbers we divide the lower bound by 2. (We need to round up after any division since the game will only have integers)

We find the following ranges are winnable and losable:

111,112 to 999,999 –> winnable
55,556 to 111,111 –> losable
6,173 to 55,555 –> winnable
3,087 to 6,172 –> losable
343 to 3,086 –> winnable
172 to 342 –> losable
20 to 171 –> winnable
10 to 19 –> losable
2 to 9 –> winnable
1 –> losable

Solution: Alice should always lose

Alice begins the game with 1 which we reasoned above was a losable number. When she presents any of the numbers 2 to 9 to Bob, he can force the total into the losing range of 10 to 18. Whatever Alice does, Bob can continue to control the resulting total so that he will win.

That’s not to say that Bob will win definitely win the game. With sloppy play, Bob can make a mistake and let Alice win.

For instance, suppose Alice starts with 9. If Bob multiplies that total by 9 to get to 72, then he has given Alice a winnable number which would allow her to control the game.

Puzzle 9: Shoot your mate

You are undercover and about to make a breakthrough with a mob boss. But your partner, not knowing your mission, tries to save you and gets captured.

The boss suspects you might be working with the authorities. He asks you to prove your allegiance. He gives you a gun and requests you shoot your partner. If you don’t fire, you will surely be caught.

Do you do it? Why or why not? Use game theory reasoning to figure it out.

Answer to Puzzle 9: Shoot your mate

Let's think about the problem strategically. The boss either trusts you, or he does not, and he has either handed you a loaded gun or not.

Imagine for a second the boss has in fact handed you a loaded gun. That would only be reasonable if he truly trusted you. Right? After all, if he handed you a loaded gun but thought you were a spy, then he would have to be worried you could fire the gun at him.

It only makes sense to give a loaded gun to someone you deeply trust. But in that case, there is no reason to test the person's loyalty!

The very fact you are being tested means the boss does not trust you. And in that case, the only sensible thing for the boss is to hand you an unloaded gun.

We can write out a matrix that shows handing you a loaded gun is a weakly dominant strategy. It is simply safer to do, whether he trusts you or not.

		Gun: Loaded	Gun: Not
Villain	Trusts you	ok	ok
	Not	risky!	ok

Therefore, if you are asked to shoot your mate, you can be reasonably sure the gun is not loaded. You should shoot at your partner to keep your cover and pray the boss was not crazy enough to hand you a loaded gun (of course, a villain as sadistic as the Joker might do this).

Examples in the show 24 (mild spoilers)

Jack Bauer is a game theorist. There are a couple of memorable instances of this trope that I want to mention here. (There are plenty of other examples in tv, movies, and literature at tvtrope.org)

Example from Season 4

In Season 4, a Muslim terrorist Dina defects to the American side to protect her son. She helps Jack Bauer to find the terrorist leader Marwan, who then questions her trust.

Marwan offers Dina a gun and tells her to shoot Jack to prove her loyalty. Dina gets nervous, because if she kills Jack then she would risk the federal protection on her son.

Dina shoots at Marwan only to find the gun is not loaded. Her deception is revealed and Marwan orders her to be shot.

Example from Season 3

Another instance happens in Season 3 when Jack was in deep cover with the Salazar brothers. Jack's partner, Chase, does not realize this and he makes a heroic effort to rescue Jack.

Unfortunately Chase is apprehended and it raises doubts whether Jack is secretly working with authorities. Ramon Salazar hands Jack a gun and tells him to shoot Chase to prove he is trustworthy. Jack takes fire, and it turns out the gun was not loaded so Chase survives. Jack keeps his cover and eventually saves the day as usual.

Chase later finds out Jack was undercover, and he is deeply angry that Jack took aim.

Jack reveals his game theoretic thinking all along, in Season 3, episode 13.

> Chase: You put a gun to my head, and you pulled the trigger.
>
> Jack: I made a judgment call that Ramon Salazar would not give me a loaded weapon–that he was testing me.
>
> Chase: And what if it is was loaded? Then what?
>
> Jack: Then I'd have finished my mission.

It is not an easy thing to shoot at your mate, but it is a judgment call that fits in line with strategic thinking.

Puzzle 10: When to fire in a duel

What's the right time to shoot in a duel?

A simple dueling game

Consider a duel between two players A and B, in which each person has one bullet.

A and B start the duel at a distance $t = 1$ from each other. They approach each other at the same speed, and each has to decide when to shoot.

As they get closer to each other, their accuracy increases. At distance t, the player A has a probability $a(t)$ of killing his opponent, and for player B it is $b(t)$. Assume both players are aware of the other's skill.

In this duel, missing your shot is very costly. If a player shoots and misses, then the other player keeps approaching until he gets to point blank range and shoots with complete accuracy.

What is the optimal strategy of this game? That is, at what point should each player shoot?

Answer to Puzzle 10: When to fire in a duel

We will separately solve for the best time for each player to shoot.

When player A should fire

The tricky part to the game is balancing when to shoot. If you fire too early, then your opponent kills you for sure. If you wait too long, then you can also get beaten if your opponent is a good shot.

We can think about when player A should fire by listing out the chance of surviving in the different possibilities of firing at point t.

> If player A fires first: Player A will survive only if he hits, which happens with probability a(t)

If player A waits to fire: Player A survives only if player B misses, which happens with probability $1 - b(t)$

Now we can reason out player A's strategy. Player A will want to fire first if his probability of hitting is greater than player B's probability of missing:

$$a(t) \geq 1 - b(t)$$

But he must also be careful not to fire too early. He should always wait if his probability of hitting is smaller than player B's probability of missing:

$$a(t) \leq 1 - b(t)$$

Putting those two equations together, we can see that Player A should shoot at the time when he is at distance t^* where

$$a(t^*) = 1 - b(t^*)$$

Or alternately written,

$$a(t^*) + b(t^*) = 1$$

Solution: when player B should fire

We can do the same exercise for player B. Notice the same conditions are true:

> If player B fires first: Player B will survive only if he hits, which happens with probability b(t)
>
> If player B waits to fire: Player B survives only if player A misses, which happens with probability 1 – a(t)

Now we can reason out player B's strategy. Player B will fire first, if his chance of hitting is better than his opponent's probability of missing:

$$b(t) \geq 1 - a(t)$$

But he must also be sure not to fire too soon. He needs to wait so long as his chance of hitting is smaller than his opponent's probability of missing:

$$b(t) \leq 1 - a(t)$$

Putting those two equations together, we can see that Player B should shoot at the time when he is distance t^* where

$$b(t^*) = 1 - a(t^*)$$

Or alternately written,

$$a(t^*) + b(t^*) = 1$$

Solution: they fire at the same time!

From the equations, you'll notice that both players find their optimal firing times satisfy the same equation. That is, they both choose to fire at the same time! There is one specific time and distance which is optimal for both players.

This would not be surprising if the two players had the same accuracy level. But we solved this game using the assumption their accuracy levels were different.

So why do they end up shooting at the same time?

We can reason why this must be the case. If one person chose to fire earlier than another, say 5 seconds earlier, then he would be better off waiting. His opponent is not shooting for another 5 seconds, so he might as well wait a few more seconds to get closer and increase his accuracy.

As the equations show above, the right time to shoot is just when your chance of hitting equals your opponents chance of missing. And since one person's failure is another person's success, this means both players choose the same time when they are a distance such that their accuracy functions sum to a probability of 1.

Puzzle 11: Cannibal game theory

A traveler gets lost on a deserted island and finds himself surrounded by a group of n cannibals.

Each cannibal wants to eat the traveler but, as each knows, there is a risk. A cannibal that attacks and eats the traveler would become tired and defenseless. After he eats, he would become an easy target for another cannibal (who would also become tired and defenseless after eating).

The cannibals are all hungry, but they cannot trust each other to cooperate. The cannibals happen to be well versed in game theory, so they will think before making a move.

Does the nearest cannibal, or any cannibal in the group, devour the lost traveler?

Answer to Puzzle 11: Cannibal game theory

I find the problem interesting because the solution uses two types of induction: the strategy depends on backwards induction, and the proof is based on mathematical induction.

The short answer is the traveler's fate depends on the parity of the group. If there is an odd number of cannibals, the traveler will be eaten, but if there is an even number, the traveler will survive.

To prove this, we will consider small groups and use mathematical induction to explain the solution for larger groups.

Case n = 1: this is an obvious case. If there is one cannibal, the traveler will be eaten. It doesn't matter that the cannibal will get tired because there are no other cannibals around as a threat.

Case n = 2: this is a more interesting case. Each cannibal wishes to each the traveler, but each knows he cannot. If either cannibal eats the traveler, then he will become defenseless and the other one will eat him. So each cannibal uses backwards induction to realize that the only strategy is to not eat the traveler. The hapless traveler finds a bit of luck, therefore, and actually survives.

Case n = 3: this is where the problem gets interesting. The best strategy is for the closest cannibal to make a move and eat the traveler. The cannibal will be defenseless after eating, but ultimately he will be safe. Why is that? The reasoning is due to induction: once the cannibal eats the traveler, the resulting situation has 2 unfed cannibals and the 1 defenseless cannibal. But as we just showed above, when there are 2 unfed cannibals, neither will make a move for fear of being eaten by the other! Thus the first cannibal to make a move will be safe as the remaining 2 cannibals block each other.

We can prove the higher cases using mathematical induction. If the number *n* is odd, then the closest cannibal can safely eat the traveler because the remaining number of unfed cannibals is even (and by induction, with an even number of unfed cannibals no one makes a move). If the number *n* is even, then no cannibal will eat the traveler, for if he did, the remaining number of cannibals would be odd, meaning he will get eaten by the induction hypothesis.

I should point out this problem highlights one of the problems with game theory and backwards induction.

If a group had 20 cannibals, the traveler would be safe. But if the group had 19, the traveler would die. That's all fine in theory, but that's a HUGE difference in outcomes for a detail like group size.

What if some of the cannibals counted wrong, or if as the action took place another cannibal appeared to throw off the rational outcome?

So unfortunately there are games that have mathematical elegance but shed light on how seemingly minor details can have profound effects on the game theory prediction.

Puzzle 12: Dollar auction game

A teacher holds up a dollar bill to a class and announces the money is up for sale.

Bidding starts at 5 cents and bids will increase by five cent increments.

There are two rules to when the game ends.

1. The auction ends when no one bids higher. The highest bidder pays the price of his bid and gets the dollar as a prize.

2. The second highest bidder is also forced to pay his losing bid (5 cents less than the winning bid) but gets nothing in return.

How will this game play out? What is your best strategy?

Answer to Puzzle 12: Dollar auction game

Like many economic students, I learned about this game first-hand. My teacher described the game as a chance for us poor students to make a small profit, if we were smart enough. Little did we know how much he was tricking us.

The bidding began at 5 cents and a hand shot up to claim the bid. Would anyone pay 10 cents? Another hand shot up.

What about 15 cents? Again, another hand shot up. Bidding at this stage seemed harmless because it's an obvious deal to buy a dollar for any amount less.

The twist became clear about when the high bid was 75 cents. Many students started to think about how the second rule–the one requiring the loser to pay–would affect incentives.

What might the second highest bidder think at this stage? He was offering 70 cents but being outbid. There were two choices he could make:

--do nothing and lose 70 cents if the auction ended

--bid up to 80 cents, and if the auction ended, win the dollar, and profit 20 cents

It's also possible a third person comes and takes the top spot, but that's not an action one can necessarily depend upon. So ignoring this option, the better choice is to bid 80 cents rather than do nothing.

But this action has an effect on the person bidding 75 cent, who is now the second highest bidder. This person will now make a similar calculation. He can either stand pat and stand to lose 75 cents if the auction ends, or he can raise the bid to 85 cents and have a chance of profiting 15 cents. Again, bidding higher makes sense. Thinking more generally, it always make sense for the second highest bidder to increase the bid.

Such strategy is why the game unraveled pretty quickly. Soon cash-strapped students were bidding more than one dollar and fighting over who would lose less money.

It is the incentives that dictate this weird outcome. Consider an example when the highest bid is $1.50. Since the high bid is above the prize of $1, it is clear no new bidder will enter. Hence, the second bidder faces the two choices of doing nothing and losing $1.45, or raising the bid to $1.55 to lose only 55 cents if the auction ends.

In this case, it makes just as much sense to limit loss as it does to seek profit. The second highest bidder will raise the bid. In turn, the other bidder will perform a similar calculation and again raise the top bid. This bidding war can theoretically continue indefinitely. In practical situations, it ends when someone chooses to fold.

In my class, the game ended around $2 when one player decided to end the madness. But talk to economics professors and you'll hear that it is not unusual for the game to end at anywhere from $5 to $10. It's especially juicy if the two bidders dislike each other in social circles, and that adds its own element of entertainment. As a side note, the game can be played in other bid increments too, like 1 cent or 10 cents.

I think the game offers two insights. First, it is best to avoid such games from the outset. And second, if you find yourself in one, cut your losses early. It is better to lose at 2 cents than at 2 dollars.

Can the auction be gamed?

At this point you might be asking if the problem is competition. Could cooperation lead to a good outcome? In theory, yes. It is possible to cut a deal with others to avoid the bidding war. Imagine a class of 9 students who wanted to embarrass a professor. One person could bid a single penny, everyone could agree not to bid higher, ending the game, and the profit of 99 cents could be shared as 11 cents per person.

The solution is promising but the problem is getting everyone to cooperate. Every person has incentive to deviate. Imagine one person who wants to show up the leader. If he bet 2 cents, he has a chance of getting 98 cents for himself rather than settling for the meager split of 11 cents.

Nothing holds players to their words, and when strangers are involved, there is really no guarantee or time to plan in advance. This is why a large lecture hall or sizeable dinner party provide suitable locations to play this game.

I've been talking about the game very negatively so far, but there is always another side to the story. Although buyers fare poorly, the auctioneer makes out like a bandit. It's a trick that makes even rational buyers overspend vast amounts. It's no wonder that economics professors love to hold this auction.

Puzzle 13: Bottle imp paradox

A stranger catches your attention one day. He offers you an interesting proposition.

He wants to sell you a bottle that contains a genie that will grant you any wish you want.

There is only one catch: you must sell the bottle at a loss within one year, or you'll be condemned to misery in Hell for the afterlife.

The stranger asks you what you would like to do.

–Do you buy the bottle?

–What price do you pay?

–What is the lowest price one should buy the bottle for?

Answer to Puzzle 13: Bottle imp paradox

You first consider the price. On the one hand, you do not want to pay too high a price. You worry about shelling out cash which you cannot recover until you sell. On the other hand, you do not want to pay too low a price, or else you risk not finding another buyer.

What price is sensible? Let's start from the beginning and work our way up. Suppose you offer to pay only one cent. This turns out to be a very bad price. The reason is there is no lower denomination and hence it will be impossible to find another buyer. You will be stuck with the bottle. Buying the bottle at one cent is equivalent to buying your own eternal damnation. So clearly this is a bad price.

But what about two cents? At first, this seems okay. If you buy at two cents, then you could theoretically sell for one cent. The problem is that you will be hard pressed to find a buyer. The reason is the person who buys from you is buying at one cent. And as argued just above, it is stupid to buy the bottle at one cent. Therefore, no one would want to buy the bottle at two cents.

Indeed, this logic can be extended. No one should want to buy at three cents, or four cents, and so on. Inductively one can reason there is no "safe"? price to buy the bottle. Thus, the bottle should never be bought because it will be hard or impossible to find a buyer.

But in practice, this conclusion feels wrong. You would expect a buyer at a high enough price. If you buy the bottle for $100, for example, you can likely find someone who will want to buy at $99.99. And they will feel safe, reasoning that they can find someone willing to pay $99.98, and so on.

The bottle imp paradox is that inductive reasoning and practical reasoning come to contradictory conclusions. Is the bottle never to be bought, or is there some high enough price range?

How can we resolve this paradox? I'll present two reasonable resolutions.

Resolution 1: the sinner saves the day

The paradox could be readily resolved with the existence of an atheist buyer. There could be someone who buys the bottle without expecting to sell it. This may be someone who does not believe in a supernatural afterlife with damnation.

Or alternately, it could be a buyer who is a sinner that cannot be saved. Since his life is already destined for damnation, having the bottle does not add an additional penalty.

The latter situation is more or less the resolution offered in Stevenson's story *The Bottle Imp,* on which this puzzle is based.

Resolution 2: foreign currencies

Another trick is that are currencies with money worth less than one penny. In Stevenson's story, the protagonist travels to Tahiti in search of a coin worth one-fifth of an American penny.

Introducing foreign currencies also allows for the bottle to be sold indefinitely. The reason is that currencies fluctuate in values on the foreign exchange market. One could buy the bottle for a low price in one currency, and sell it when the currency appreciates. The bottle could be sold back and forth in accordance with the swings of the market.

Of course, now we are back to the situation of betting on the market and the madness of men, which is not all that comforting.

Puzzle 14: Guess the number

On a game show, two people are assigned whole, positive numbers. Secretly each is told his number and that the two numbers are consecutive. The point of the game is to guess the other number.

Here are the rules of the game:

–The two sit in a room which has a clock that strikes every minute on the minute

–The players cannot communicate in any way

–The two wait in the room until someone knows the other person's number. At that point, the person waits until the next strike of the clock and can announce the numbers

–The game continues indefinitely until someone makes a guess

–The contestants win $1 million if correct, and nothing if they are wrong

Can they win this game? If so, how?

(Credit: I came across a highly amusing puzzle in Impossible?: Surprising Solutions to Counterintuitive Conundrums)

Answer to Puzzle 14: Guess the number

At first it seems like the contestants can do no better than chance. If a contestant has the number 20, for instance, there is no way to tell if the other person has 19 or 21.

The naive strategy is to wait a bit and then take a guess at the other number, yielding a 50 percent chance of success.

But can they do better?

The solution

The answer lies in the subtle rule about announcing the number once the clock strikes. It turns out that two players who reason correctly can win every single time, if they think inductively.

To see why this is so, think about a case where the players can win for sure. If one of the players gets the number 1, then he can be sure the other player has the number 2. There is no other possibility because the two assigned numbers are consecutive and positive. Therefore, this player will immediately deduce the answer and announce the numbers during the first strike of the clock.

Now consider when the players are assigned the numbers 2 and 3. The player with 2 can reason as follows. "I know my partner can either have 1 or 3. But if he has 1, then he should know it and will announce it at the first strike of the clock. So if the clock strikes once and nothing happens, I can be sure that I have the lower number. Therefore our numbers must be 2 and 3." So what will happen is this player will announce the numbers at the second strike of the clock.

This reasoning can naturally be extended by induction. The general statement is the player with the number *k* will realize the other has $k + 1$ and can announce this information at the k^{th} strike of the clock.

They can win every time!

Final thought: the connection with common knowledge

The puzzle illustrates the game theory concept of common knowledge which is distinguished from the weaker concept of mutual knowledge.

Roughly speaking, an event is mutual knowledge if all players know it. Common knowledge also requires that all players know the event, all players know that all players know it, and so on ad infinitum.

Here is how the two concepts work in the game. When a player has the number 20, it is mutual knowledge that neither player has the number 1, or 2, or so on up to 18. But that deduction is not good enough to solve the game.

And that is where the clock provides a helping hand. When the clock strikes once, and no one answers, the fact that neither player has 1 transforms from mutual knowledge into common knowledge. This alone is trivial given the numbers are consecutive, but as shown above, this stronger set of knowledge can build up on consequent strikes of the clock. Each time the clock strikes the set of excluded numbers is included as common knowledge and eventually the players can win.

Puzzle 15: Guess 2/3 of the average

A teacher announces a game to the class. Here are the rules.

1. Everyone secretly submits a whole number from 0 to 20.

2. All entries will be collected, and the guesses will be averaged together.

3. The winning number will be chosen as two-thirds of the average, rounded to the closest number. For instance, if the average of all entries was 3, then the winning number would be chosen as 2. Or if the average was 4, the winning number would be 3 (rounded from 2.6666…).

4. Entries closest to the winning number get a prize of meeting with the professor over a $5 smoothie. (In the academic version of the game, multiple winners split the prize, but this teacher is being generous).

What guess would you make? What if everyone were rational?

Answer to Puzzle 15: Guess 2/3 of the average

The game is called a "*p*-beauty contest." The "*p*" refers to the proportion the average is multiplied by–in this case, *p* is two-thirds. If you're wondering, the game has the same flavor for any value of *p* less than 1. Why is it called a beauty contest? It's because the game is the numbers-analog to a beauty contest developed by John Maynard Keynes.

Here is the beauty contest that Keynes pondered. Imagine that a newspaper runs a contest to determine the prettiest face in town. Readers can vote for the prettiest face, and the face with the most votes will be the winner. Readers voting for the prettiest face will be entered in a raffle for a big prize.

How does the game play out? Keynes wanted to point out the group dynamics. The naive strategy would be to pick the face you found to be the most attractive. A better would be to picking the face that you think *other* people will find attractive.

The number "beauty contest" has the same kind of logic. You don't pick a number you like. You pick a number that's closest to two-thirds of the average of everyone else. The twist of both games is that your guess affects the average outcome. And each person is trying to outsmart everyone else.

My experience with the game

The puzzle is based on my own experience in a game theory course.

Given the subtlety of the game, my professor was banking on paying out to only a few winners. Although it was mathematically possible for each of us to win, and he was taking that risk. In fact, he knew that if we were all rational, we would all win. He would have to pay out a $5 smoothie to 50 students–that is, he made a $250 gamble playing this game.

Why was he so confident? Let's explore the solution to the game and see why it's hard to be rational.

Numbers You Shouldn't Pick

Even though it's not possible to know what other people are guessing, this game has a solution. If everyone acts rationally, there are only two possible winning numbers. It takes some crafty thinking, but it is really based on two principles I think you will accept.

Principle 1: Don't Play Stupid Strategies (Eliminate Dominated Strategies)

The first principle is that players should avoid writing down numbers that could never win. That sounds logical enough, but it's not always the case. We all can agree that writing a number that could never win is just a dumb, stupid strategy. You are picking an option that's inferior to something else, and hence is known as a dominated strategy.

Are there any dominated strategies in the beauty contest?

To start answering that question, we need to figure out which numbers will never win. A natural question is what is the highest winning number? You would never want to pick a number larger than that, unless you want to lose.

You know that the highest number anyone can pick is the number 20. If every single person picked 20, then the average would be

20. The winning number would be two-thirds of 20, which is 13 when rounded.

Should you ever find yourself submitting 20?

The answer is no–there is always a better choice, say the number 19. The only time 20 wins is precisely when everyone else picks it and everyone shares the prize. In that case, you would be better off writing 19 to win the prize unshared. Plus, by writing 19, you can possibly win in other cases, like when everyone picks 19. You are always better off writing 19 than 20. The guess of 20 is dominated–it's dumb.

You should never choose 20. And your rational opponents should be thinking the same way. So here's the big result: you can reason that no player should ever choose 20.

Principle 2: Trim the Game, and Apply Principle 1 Again (Iterate the Elimination Process)

By principle 1, no player will ever choose 20. Therefore, you can essentially remove 20 as a choice. The game trims to a smaller beauty contest in which everyone is picking a number between 0 and 19. The smaller game has survived one round of principle 1.

Now, repeat! Ask yourself: in the reduced game, are there any dominated strategies?

Now 19 takes the role of 20 from the last analysis. Since 19 is the highest possible average, it will never be a good idea to guess it. Applying principle 1, you can reason that 18 is always a better choice than 19. Thus, 19 is dominated and should be eliminated as a choice for every player.

The game is now trimmed to picking numbers from 0 to 18. This is the result of two iterations of principle 1.

There's no reason to stop now. You can iterate principle 1 to successively eliminate choices of 18, 17, 16, and so on. The only numbers remaining will be 0 and 1. (This requires 19 iterations of principle 1.)

There is a name for this thought process. It's aptly named, but a mouthful: iterated elimination of dominated strategies (IEDS). The idea is to eliminate bad moves, trim the game, and iterate the process to find the surviving moves.

These remaining strategies are considered to be rationalizable moves, that is, moves that can possibly win.

The Equilibria

The only reasonable choices are to pick the numbers 0 and 1. Is either a better choice? This is unfortunately where IEDS cannot give insight.

It's possible to have 0 as a winning number–imagine all players picked 0. (The winning number will be 0).

It is also possible to have 1 as a winning number–imagine all players picked 1. (The winning number is 2/3, which rounds to 1.)

The answer will depend on what people think others will be guessing. Both equilibriums–all 0 and all 1–are achievable.

Back to the Classroom

None of us in the class had this deep understanding of IEDS. We were just learning game theory–it was actually our third lecture. My professor was pretty sure our guesses would be all over the place.

But Stanford kids can be crafty. One student used some sharp thinking and realized that coordination would help; he asked if we could talk to each other. The professor, still feeling we were novices, confidently replied with a smile, "Sure. Go ahead." We only had 10 seconds to write down our answers anyway.

Before the professor could change his mind, the student quickly shouted to all of us, "If we all write down 0, we all win."

It was remarkable. He figured out the equilibrium and told us what to do! He couldn't be tricking us because the math was clear: if we all picked 0, we would all have winning numbers.

My professor's face seemed to drop. That's $250 on the line. (He never let future years talk before their bids).

How Smart Are Stanford Kids?

The professor was relieved after he tallied the votes. He told us that admirably most of us wrote down the number 0 (I was among those who did). But there were larger answers too, ranging from 1 to 10.

Someone actually wrote down 10! And this was after being told the answer.

After all was said and done, the winning number turned out to be 2, and the prize was awarded to three students. Thanks to our irrationality, my professor only paid out $15.

It was even better. My professor grilled the students who wrote down larger numbers. They all squirmed, as he was physically intimidating, and explained reasons like "it was my lucky number" or "I don't know. I wasn't really thinking."

The Practical Lesson

What is going on? This is a group of smart students that was told the answer to the game.

The example illustrates a flaw of IEDS. It can get you reasonable answers if you think players are reasoning out further and further in nested logic. We don't have infinite rationality, only bounded rationality.

The practical answer to what you should write depends on the book answer plus your subjective beliefs about what other people do. It's the combination of book smarts plus social smarts that matters.

The people who wrote down the winning numbers told the class they suspected some people would deviate for irrational reasons. And they were rewarded for not confusing theory and practice.

Puzzle 16: Number elimination game

Bored at the airport, Alice and Bob decide to play the following mathematical game.

Alice writes the numbers 1, 2, . . . , N on a piece of paper.

Bob goes first, and he picks two numbers x and y from the list. Bob crosses out these numbers from the sequence, and he includes a new number equal to their positive difference (in other words, he puts $|x - y|$ on the list).

Alice takes her turn and does the exact same thing with the remaining numbers on the list.

Bob and Alice continue to play, in turn, until there is only one number left.

Alice wins if the final number is odd, and Bob if even.

1. What strategy should Alice and Bob have for the game?

2. Is there a winning strategy?

Answer to Puzzle 16: Number elimination game

Let's work through an example to see how the game might play out.

Suppose Alice just writes the numbers 1, 2, and 3 in the initial list.

Bob can choose any two numbers, which means he can make any of the following three moves:

> –if he chooses (1, 2) the resulting list is 1, 3
>
> –if he chooses (1, 3) the resulting list is 2, 2
>
> –if he chooses (2, 3) the resulting list is 1, 1

Alice will simply have to pick the two numbers that are remaining, and the results will be as follows:

–if the list was 1, 3 the resulting number is 2 (which is even so Bob wins)

–if the list was 2, 2 the resulting number is 0 (which is even so Bob wins)

–if the list was 1, 1 the resulting number is 0 (again an even number, so Bob wins)

This worked example shows Alice will lose the game regardless of how she or Bob choose to play.

But why is that? And are there games that Alice can win?

The trick to solving the game

The trick is to notice what is happening on each turn of the game.

In the version with the numbers 1, 2, 3 worked out above, we can notice something interesting about the sum of the numbers in the list.

The original sum of all the numbers is 6, an even number.

When Bob moves, the resulting sums can either be 4–if he picks (1, 2) or (1, 3)–or the sum can be 2–if he picks (2, 3). In either case, the sum of the numbers is even.

And after Alice moves, we showed the resulting number must be even as well, which is why she loses.

This suggests a pattern: the final number, as well as every intermediate sum, will have the same *parity* (the property of being odd or even) as the sum of the original list.

If true, that means Alice wins if the original list has an odd sum, and Bob wins if the original list has an even sum.

But how can we prove this is true?

Proof that parity is unchanged / invariant

Suppose the original sum of the numbers is labeled S.

On Bob's turn, he removes two numbers $x > y$ from the list, and he writes another number $(x - y)$.

This means Bob's action reduces the original sum S by the following:

$$x + y - (x - y) = 2y$$

The thing to notice is that $2y$ is an even number, which means the parity of the sum is unchanged by a move in the game. In other words:

–If the original sum S was even, then each turn it is reduced by an even number. As an even minus an even is also an even number, this means every intermediate sum will be an even number. Hence, the final number must be even as well.

–If the original sum S was odd, then each turn it is reduced by an even number. As an odd minus an even is an odd number, this means every intermediate sum will be an odd number. Hence, the final number must be odd as well.

Answering the original questions

1. What strategy should Alice and Bob have for the game?

This is something of a trick question.

Alice and Bob have no particular strategy for play: the game is decided by the parity of the sum of the initial list.

Note that if the initial list is 1, 2, 3, . . ., N, then its sum is $N(N+1)/2$. The parity of this number decides who wins the game.

2. Is there a winning strategy?

Again this is a trick question as the moves are irrelevant to the outcome.

Arguably, the winning strategy is to influence the choice of the initial list, and hope the other person doesn't notice.

Puzzle 17: Hat puzzle

There are three players in this game. Each player has either a blue or red hat placed on their head, determined randomly.

Each player can see the colors on the other two player's hats, but not the color of their own hat.

Each person in the group must simultaneously write the color of their own hat on a piece of paper, or they can write "pass". The group loses if someone writes a wrong guess, or if they all write "pass."

No communication is allowed, except for an initial strategy session.

What strategy should the group use, and what is their chance of success?

Answer to Puzzle 17: Hat puzzle

The Basic Strategy (50 percent winning)

The rules heavily penalize incorrect guesses. A single incorrect guess makes the group lose–even if the other two players guess correctly. A single incorrect guess is the apple that spoils the bunch.

So it's important the rules allow for players to pass. If a player doesn't have a good guess, it would be a good idea to pass.

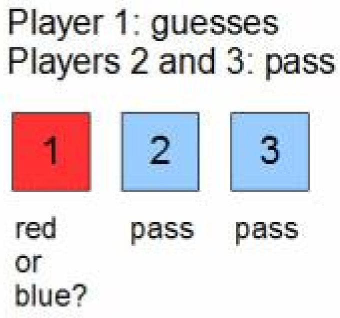

A basic strategy would be to minimize the risk of bad guesses. Force two players to pass in every game and make one person the official guesser. The group wins exactly when this person guesses correctly.

How often will the group succeed? Since the hat color is chosen by a coin flip, there is a 50 percent chance of guessing the correct color.

But can the team do better than random chance?

The trick is figuring out the players do have a way of coordinating as a group. Doing this, they can make winning an amazing 75 percent chance. Let's investigate why.

One Optimal Strategy (75 percent winning)

Motivating question: does observing the other two hat colors tell you anything about your own hat color? In other words, if you see two red hats, does that make your hat more likely to be blue?

The answer is no, and that's a potential roadblock. Regardless of what you see, your hat color is determined by a coin flip. Fair coins are never "due" for a particular outcome–each toss is independent.

But don't get caught up in probability–the fact is that seeing the other hat colors does convey information. The problem is the figuring out how to transmit that information to the other players.

To get around that, players need to coordinate guesses based on what they see. If possible, they still want to minimize bad guesses by having two people pass and one person guess. What's needed now is a group strategy.

How can they do that? It starts by taking a step back and considering the possible distributions of hat outcomes. With three players and two hat colors, there are a total of eight equally likely outcomes:

> RRR, RRB, RBR, RBB, BBB, BRR, BRB, BBR

Is there anything special about the distribution?

One feature is that most outcomes–six of them–include at least one hat of both colors. Only two extreme outcomes don't–the ones with all red hats or all blue hats.

We can analyze further. Among outcomes with both hat colors, there logically has to be two hats of one color (the "majority" color) and one hat of another color (the "minority" color).

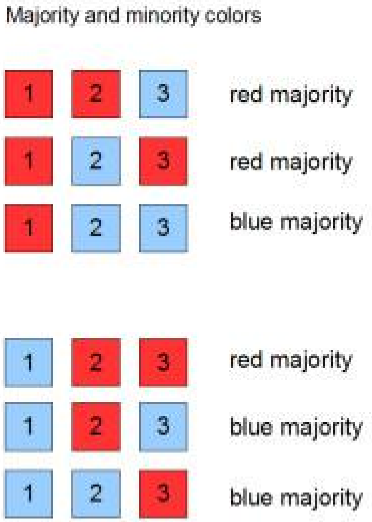

Here's the kicker: by looking at the other hats, players can identify whether they are wearing a majority color or a minority color.

For instance, if a player sees both a red and blue hat, then the player must be wearing the majority color (which could be red or blue).

If a player sees two blue or two red hats, then the player must be wearing the minority color, which will be the opposite color of what the player sees.

Now the idea is to get the player with the minority hat color to guess and force the other people to pass.

So here is the strategy:

> if you see both a red and a blue hat, then "pass"
>
> if you see two red hats, then guess "blue"
>
> if you see two blue hats, then guess "red"

This strategy wins in all six cases with at least one hat of each color. It only loses in the two cases of all-red or all-blue, in which all players guess incorrectly.

Here is how players would guess:

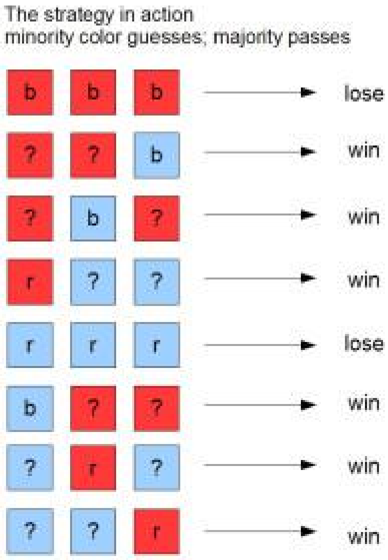

All told, the group wins in six of eight possible outcomes–a whopping 75 percent chance.

Extension: The host can learn

If you're playing rock-paper-scissors against a computer that mixes randomly, you could win 1/3 of the time simply by picking one strategy, say rock. But if the computer could learn and analyze your pattern, it might respond by countering with paper and start winning a lot. To maintain your 1/3 winning chance, you need to randomize your choices among rock, paper, and scissors.

In the hat game, the players have a 75 percent chance of winning, but the strategy has a pattern. It loses every time the hat colors are all the same. A responsive host, like the computer in rock-paper-scissors, would see the pattern and respond by assigning hats to be the all one color more frequently. To keep the host honest, the players need to randomize.

Is there some a way the players can win, without creating a pattern of outcomes in which they all lose?

Amazingly, yes there is! Even more surprising, the winning percentage stays at 75 percent.

The Random Optimal Strategy (75 percent winning)

The random strategy is a refinement on the static one given above. The key to the above strategy is that players essentially bet against the outcome being all-red or all-blue. Knowing that, it was possible to coordinate guesses so only one person guessed and gave a correct answer.

There's nothing special about picking all-red or all-blue.

The players can randomly pick any color combination and its "opposite" configuration (red-blue-red and blue-red-blue are opposites) as outcomes to bet against. The remaining six outcomes can be coded based on the hat colors that each player sees.

Why would this work, and why does it have to be "opposite" combinations?

The eight outcomes of the hat game can be visualized as vertices of a cube. Adjacent vertices differ by changing only a single "coordinate," that is, the color of one player's hat.

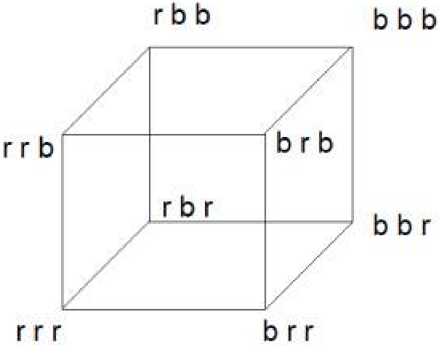

The graphical interpretation is as follows: can the players identify which vertex they belong to? The information they are given is the other two hat colors they see–that is, they are effectively placed at midway points along the adjacent edges.

Each player can see the coordinates of the other two players, but is unsure about the own coordinate–that is, the player is unsure which of the two possible endpoints the group belongs to.

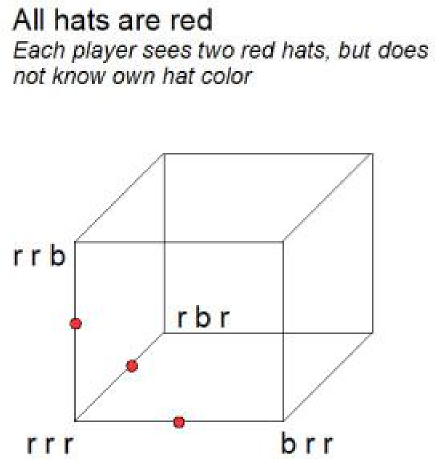

We want a situation where exactly two players will not be able to tell the vertex (they will “pass”) and the remaining player will know the location (and guess correctly).

Such a unique coding occurs when players bet against a random vertex and the “opposite” one–a limitation that gives maximal location information.

In any of these cases, players only lose if in fact the outcome is one of the two they bet against, meaning they have a 75 percent chance of winning.

The 75 percent chance of winning applies to every time the game is played but the losing outcomes are randomized. Hence, the host won't be able to exploit any particular color combination.

Appendix: The coding for betting against red-blue-red and blue-red-blue outcomes (others can be found similarly)

Strategies:

Player 1:
If see blue-red, then pick red.
If see red-blue, then pick blue.
Else pass.

Player 2:

If see red-red, then pick blue.
If see blue-blue, then pick red.
Else pass.

Player 3:
If see red-blue, then pick red.
If see blue-red, then pick blue
Else pass.

Puzzle 18: Polynomial guessing game

Alice and Bob decide to play a math game. Alice secretly writes down any polynomial p(x) of one variable that she wants. The polynomial can be of any degree, but to limit the scope somewhat, the polynomial can only have nonnegative integer coefficients.

Thus, Alice can pick polynomials like $2x^2 + 1$ or $3x^{100} + 2x^2$, but she cannot pick polynomials with negative coefficients like $-x^2 + x$, or non-integer coefficients such as $0.5x^2$.

Bob has to guess the polynomial. He gets to ask two questions of Alice. First, he gets to pick any number *a* and ask Alice for the value of p(*a*). Then, he gets to pick another number *b* and ask for the value of p(*b*).

Bob wins if he can guess the polynomial; otherwise Alice wins.

After playing the game for several rounds, Bob announces that he has a winning strategy. Can you figure out what it is?

(Credit: the puzzle was originally submitted as “A Perplexing Polynomial Puzzle” in College Mathematics Journal, March 2005, p. 100.)

Answer to Puzzle 18: Polynomial guessing game

Before I explain the answer, let's see the strategy in action.

Alice picks a polynomial, and then Bob asks for the value of p(1). Let's say that Alice replied p(1) = 4.

Bob then asks for the value of p(5). Alice replies the answer is 36.

Bob thinks for a moment, and then announces polynomial must be $x^2 + 2x + 1$. And he's right!

How did Bob do this?

Rather than explain the answer right away, I want to do one more example.

Alice comes up with another polynomial and Bob again picks the number 1. Alice replies that p(1) = 9.

Bob then picks the number 10, and he learns that p(10) = 432.

Suddenly Bob exclaims the answer must be $4x^2 + 3x + 2$, and he's right again!

This example is the secret to the whole puzzle. The interesting part is that Bob's second question led to the answer 432, and the polynomial had coefficients of 4,3, and 2 for its descending powers.

Why did this happen?

The short answer is this. What Bob is doing is: he is first learning the value to p(1), and then he asks for the value of p(1)+1. It turns out every time that the digits of this answer, p(p(1)+1) in the numerical base of p(1)+1 are precisely the coefficients of the polynomial.

Why is that? Let's think critically about what is happening in each step.

The polynomial p(x) can be written as $p(x) = a_n x^n + a_{n-1} x^{n-1} + ... + a_0$

When Bob asks for the value of p(1), he will end up with the sum of the coefficients because $p(1) = a_n + a_{n-1} + ... + a_0$

Now comes the neat trick. When Bob asks for the value of p(p+1), he ends up with the following term:

$$p(p+1) = a_n (p+1)^n + a_{n-1} (p+1)^{n-1} + ... + a_0$$

Notice anything interesting about this series?

The trick is this: each coefficient is uniquely attached to a different power of p + 1. By construction, each coefficient is smaller than the attached term p + 1. Therefore, the series is a unique representation of the number p(p + 1) in the number base p + 1!

That is, if we write p(p+1) in the number base p + 1 we get the number:

$a_n \, a_{n-1} \ldots a_0$ (base p + 1)

This representation is not possible if you have negative or non-integer coefficients, and hence the restriction.

(Notice we could equivalently use any number larger than p + 1, or simply any number larger than the maximum coefficient. But p + 1 is the smallest value guaranteed to work)

In the example above, when we found p(10) was 432, we could view the number in base

10 as the coefficients of the polynomial. In other words, we could deduce the polynomial must have been of degree 2, and the coefficients of the polynomial were 4, 3 and 2 in descending order.

In the other example when p(5) was 36, we needed to do a little more work. We needed to take the number 36 in base 5. This turns out to be 121: $1*5^2 + 2*5^1 + 1*5^0$. Thus we could deduce the polynomial had coefficients 1, 2, and 1.

So to summarize, Bob's winning strategy is this:

1. Ask for the value of p(1)
2. Ask for the value of p(p(1)+1)
3. Convert the value into base p(1)+1
4. The digits of the number are the coefficients of the polynomial in descending order

It's quite a remarkable and ingenious strategy.

Puzzle 19: Chances of meeting a friend

On a Friday night, two friends agree to meet up in a bar between midnight and 1 am. Each forgets the exact time they are supposed to meet, so each shows up at a random time.

Suppose that after arriving randomly, each waits 10 minutes for the other person before leaving. What is the chance the two will meet at the bar?

Game theory extension

If both friends are rational, and they want to maximize the chance of meeting the other, what strategy should each pursue? If they play optimally, what is the chance they will meet each other? (Assume each person is aware the other will wait 10 minutes before leaving.)

Answer to Puzzle 19: Chances of meeting a friend

I will present a few solutions to the problem.

Solution part 1: geometric probability

I feel this is the most elegant way to solve the problem.

If we let x denote the time one person arrives at the bar, and y the other, we can use the notation (x, y) to denote the time in minutes, after midnight, that each person arrives at the bar.

The trick now is that we can model the situation geometrically in the Cartesian plane. The x-axis can be labeled from the time 0 minutes until 60 minutes, as can the y-axis.

Now any coordinate in the 60×60 square represents a time that the pair arrives at the bar. The coordinate (0,9) means one person shows up at midnight, the second at 12:09, and clearly they will meet because the times are within 10 minutes of each other. The coordinate (1, 51), on the other hand, corresponds to one showing up at 12:01 and the other at 12:51, a time the two will not meet.

What is the set of coordinates for which the two will meet?

The notation affords a very simple way to describe the event. We want the two coordinates to be within 10 minutes of each other. We either want the x-coordinate to be 10 units smaller than y, or 10 units bigger than y. The succinct way of writing that is we want all coordinates such that $|x - y| \leq 10$.

We will draw the lines $y = x - 10$ and $y = x + 10$ and shade the area in between these two lines to denote the event.

The resulting figure is as follows:

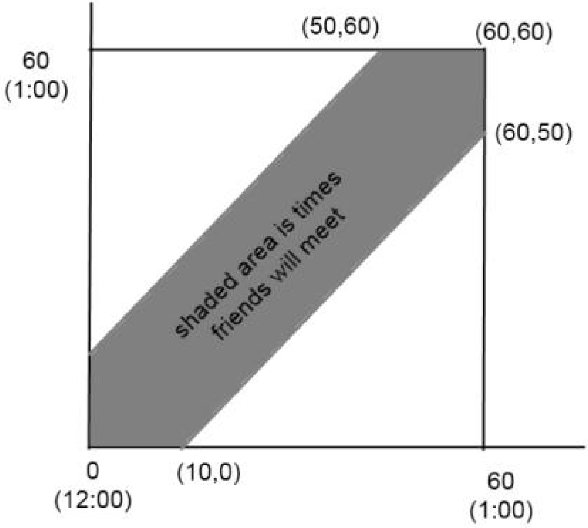

The probability of meeting is precisely the ratio of the shaded area to the total square.

This is fairly easy to calculate. The big square is 60×60 = 3,600 in area. Rather than finding the shaded area, let us calculate the unshaded area and subtract. The unshaded area consists of two isosceles right triangles with sides of 50. This means each triangle has an area of (0.5) x (50×50) = 1,250 and the total unshaded area is double that, 2,500.

The shaded area is found by subtracting the unshaded area from the total : 1,100 = 3,600 – 2,500.

Thus we conclude the chance the friends will meet is 1,100 / 3,600 =11/36, or about 30.6 percent.

Having nearly a one in three chance to meet is actually not that bad!

I find the geometric solution to be the most elegant, but it should not surprise you there are other ways to solve this puzzle.

Later in the article I will explain two other solution methods.

For now, I want to highlight another interesting fact. The math shows that even two mindless friends have a rather good chance of meeting up with each other.

In the game theory extension, I asked how likely it would be if the friends were completely rational and reasoned carefully. The surprising thing is the friends, if they reason carefully, are guaranteed to meet!

Here is why.

Game theory solution

I credit my aunt for offering a strategic answer to this mathematical problem, inspiring this extension.

In the real world, people don't show up randomly. They will use some heuristics and reason out a strategy to increase the chances of meeting their friend.

I've asked this puzzle to many people, and their reactions are quite interesting. The first thing that people notice is that it's a bad idea to show up too early or too late. Why is that?

You probably figured this logic out when solving the puzzle. If you show up right at midnight, you will only win if your friend ends up showing up after you. If you show up somewhere in the middle, you win if your friend shows up 10 minutes or less after you AND if your friend had happened to show up 10 minutes or less before you.

Similarly, you can reason it's a bad idea to show up at 1 am or near the end, in which case you win only if your friend showed up before you.

So which times are not good strategies? Let's be specific and list the out.

Dominated strategies

One time that is very stupid to show up is right at midnight. You will only win if your friend shows up during the first ten minutes–we can denote this as the interval [0, 10] for shorthand.

Rather than showing up right at midnight, you would do better to show up at 12:01. In this case you still win if your friend shows up at midnight, as he will be waiting for you. But you will also win if he shows up any time before 12:11. In short, that means you win if he shows up during the first 11 minutes of the night–corresponding to the interval [0,11].

Notice that by picking 12:01 instead of 12:00 exactly, you've increased your chances of meeting without sacrificing anything–you get an extra minute of time you will meet.

This argument shows that 12:00 is not a good strategy–it performs worse than 12:01 REGARDLESS of when the other person shows up.

We casually say arriving at 12:00 is a stupid idea. In game theory jargon, such a stupid strategy is dubbed as a dominated strategy.

(Note: it is vital that each person is aware the other person will wait 10 minutes before leaving–the rules of the game should be common knowledge)

Elimination of dominated strategies

Obviously dominated strategies should never be played. They perform worse than some other strategy, and hence they can be removed from consideration.

The above logic showed that 12:00 was dominated by 12:01. We can similarly show that 12:01 is dominated by 12:02, and in fact we can ultimately prove that showing up any time before 12:10 is a bad idea.

By exactly symmetric reasoning, we can prove that showing up any time after 12:50 is a dominated strategy. You are better off coming just a bit earlier to increase your odds of meeting up.

So that leaves us with the 40 minute interval from 12:10 to 12:50 in which both players could arrive.

If we assume the players show up randomly in this interval, then we can use a geometric argument as in Solution 1 above to find the probability of meeting jumps up to 7/16 = 43.75 percent.

But can the players do even better?

In fact they can! Here is why.

ITERATED elimination of dominated strategies

Surely the friends have reasoned this far, they are not going to stop thinking now.

We can continue to apply the same logic as before to try to trim the scope of good strategies. I mentioned this idea in the "guess 2/3 the average" puzzle.

But I will go through the reasoning again for the sake of completeness.

Remember we argued that 12:00 was a bad time to show up because it was the earliest possible time. So we concluded it was never a good idea to show up before 12:10.

That means in this reduced game that 12:10 is now the earliest time either friend would ever show up. Both friends should realize this, and we can repeat or iterate the logic again! The logical process is known as the mouthful iterated elimination of dominated strategies.

Basically 12:10 in this reduced game is very similar to 12:00 in the original game. Since no person shows up before this time, you only win if a friend shows up after you. It would make more sense to choose 12:11 to increase your chances of meeting your friend.

You are probably getting the idea, so I'll skip a few steps and get to the end result.

In the reduced game, we can prove that showing up any time before 12:20 is a bad idea. Similarly, we can prove that showing up any time after 12:40 is a bad idea.

By iterating the process of removing bad strategies, we have derived a smaller strategy space and come up with a sharper prediction of play.

The solution: iterate one more time

Notice we are down to a 20 minute interval of time from 12:20 until 12:40. There's no reason to stop here–let's iterate the decision process one more time to see if we can get any better.

In this reduced game, the earliest time one of the players will show up is 12:20. Again, we can demonstrate it's a bad idea to show up too close to the starting point of the interval.

Using the same logic as before, we can see it is best to show up only at 12:30 or later.

Using similar logic, the latest time a friend will show up is 12:40. You can see what's coming here: we can reason that it's never a good idea to show up at 12:30 or later.

Putting these two facts together, we end up with a remarkable conclusion: 12:30 is the unique arrival time (i.e. Nash equilibrium) that the friends will show up!

This solution is absolutely marvelous to me, and it even has a few other interesting properties:

> –this is an efficient time, as wait time is ZERO

> –showing up at the middle is an obvious point (known in game theory as a focal or Schelling point)

–the friends are GUARANTEED to meet up: probability of meeting is 100 percent

So two friends who are reasonable enough to think can figure out how to meet for sure without relying at all on cell phones. You can see why the world of game theory is so seductively attractive to thinking people.

I feel the fact that people cannot arrive at similar success in the real world says something about the human condition.

But anyway, let me get to some other mathematical solutions to the non-game theory version. I find these are also very satisfying.

(Extra credit: The logic is not only for 10 minutes. This would also apply for 1 minute. In fact, you can show that if each person waits for any time $t > 0$, then the unique equilibrium strategy is arrive at 12:30, leading to the outcome both people meet)

Method 2: Conditional probability

My uncle came upon this analytic solution. For narration sake, I will give a less than formal explanation–I am sure you can fill in the details.

Let's consider the game from one friend's perspective. We know the other player can show up at any time on the interval (remember in the non-game theory version any time is possible).

How likely are we to meet the other player? We can split up the cases in terms of conditional probability.

> **Case 1:** If the other person shows up in the first 10 minutes (10/60 = 1/6 of the time), the average time of showing up is 12:05. We will meet if I show up any time from 12:00 to 12:15. This interval is 15 minutes out of 60, or 1/4 of the time.

Case 2: Similarly, if the other person shows up in the last 10 minutes (10/60 = 1/6 of the time), the average time of showing up is 12:55. We will meet if I show up any time from 12:45 to 1:00. This interval is 15 minutes out of 60, or 1/4 of the time.

Case 3: Finally, if the person shows up in the middle 40 minutes (40/60 = 2/3 of the time), the average time of showing up is 12:30. We will meet if I show up any time from 12:20 until 12:40. This interval is 20 minutes out of 60, or 1/3 of the time.

These three cases cover the various conditional events. We can thus compute the probability of meeting as:

Pr(meeting) = Pr(Case 1) * Pr(meeting in Case 1) + Pr(Case 2) * Pr(meeting in

Case 2) + Pr(Case 3) * Pr(meeting in Case 3)

Pr(meeting) = (1/6)(1/4) + (2/3)(1/3) + (1/6)(1/4) = 11/36

Again, we arrive at the same solution of 11/36.

Method 3: Combinatorics

This is a solution I came upon when I was considering the practical implications of the geometric solution above.

I loved how elegant the geometric solution was, but the fact that time had to be continuous was something odd to me. I mean it is possible to show up at 12:33 and 1.14159 seconds, but who keeps track of time that accurately? And would you really be able to leave exactly 10 minutes later at 12:43 with 1.14159 seconds?

No, in the real world we are going to do some rounding, probably to the level of minutes.

So I set up a combinatorial problem as follows. Suppose each player can pick one of the minutes to arrive randomly from 0, 1, 2, ..., 60, and each person waits 10 minutes for the other person. What is the chance they will meet then?

This is a discrete version of the continuous geometric problem, so let's solve it.

Solution to discrete problem in minutes

We can proceed simply by counting the number of pairs (x, y) such that $|x - y| \le 10$ as in the continuous case.

For simplicity, let's consider the perspective of the person showing up first and count the number of integers the other person can arrive after. This will count half the cases, and we can double the result to count all cases.

–If the first person picks 0, then the person arriving second can pick times 0, 1, 2, ..., 10–there are 11 times corresponding to 0

–If one person picks 1, then the person arriving second can pick times 1, 2, ..., 11–there are 11 times corresponding to 1

–If the first person picks anything from 2 to 50, there will be 11 possible times for the person arriving second

–If the first person picks 51, the person arriving second can pick 51, 52, ... 60, or 10 cases

–If the first person picks 52, the person arriving second can pick 52, 53, ..., 60, or 9 cases

–This pattern will continue so we have 53 having 8 cases, 54 having 7 cases and so on.

To summarize, for the person arriving first, there are this many ways for the person arriving second to choose:

–For the numbers 0 to 50, there will be 11 cases

–For the numbers 51 to 60, there will be 10, 9, 8, 7, ... 1 cases, respectively

We need to double this to find the total number of winning pairs. Thus we have:

Number of times friends meet = 2[(50)(11) + 10 + 9 + 8 + ... + 1] = 2(616) = 1,232

This has to be divided by the total number of pairs. As each person can pick among 61 numbers, the total number of pairs is 61 x 61 = 3,721.

Thus the probability the friends meet in the discrete case is 33.1 percent = 1,232 / 3721.

This is actually pretty close to the continuous case. Could there be a relation between the two problems?

Solution to discrete problem in arbitrary interval

I got to thinking, what would happen if we instead split up the interval into finer points, like into seconds or so on?

I did the calculation for seconds, and I will spare you the details, but it ends up at roughly 30.6 percent. This is very, very close to the answer of the continuous model.

What would happen if the interval was divided into N pieces? And what would happen if we let N go to infinity?

If the interval was divided up as 0, 1, 2, … , N, then we need to remember that 10 minutes corresponds to 1/6 of the total time, which means it will translate into $(1/6) N$ intervals (for simplicity, let N be a multiple of 6).

Using the same logic as in the discrete case of minutes, we can count the number of ways the person arriving second could meet the person arriving first. It is:

–For the numbers 0 to $(5/6)N$, there will be $(1 + (1/6)N)$ integers of times for the person arriving second so they meet

–For the numbers $(5/6)N + 1$ to N, there will be $(1/6)N$, $(1/6)N - 1$, …. 1 times, respectively

Again, we need to double this number to account for the symmetric case. This means we have in all:

$$2\left([(5/6)N+1][(1/6)N+1)]] + \sum_{i=1}^{\frac{N}{6}} i\right)$$
$$= (11/36)N^2 + (13/6)N + 2$$

We need to divide this by the total number of cases which is $(N+1)^2 = N^2 + 2N + 1$. So the limiting case is:

$$\lim_{N\to\infty} \frac{(11/36)N^2 + (13/6)N + 2}{N^2 + 2N + 1} = \frac{11}{36}$$

When we take the limit as *N* goes to infinity this is 11/36, just as in the continuous case!

For some people it will just seem "obvious" that the discrete problem converges in limit to the continuous problem.

But anyone who has studied financial models knows that discrete versions are not the same and may not converge to the same as continuous models.

So this is an interesting result, and it's amazing how many different ways this puzzle can be solved.

Puzzle 20: Finding the right number of bidders

Alice wants to auction off a rare collector's item. She knows the item is worth somewhere between $500 and $1,000, but she has had trouble finding interested buyers.

A company offers to find interested participants at the rate of $10 per bidder. (So they'll find one bidder for $10, and ten bidders for $100)

How many bidders should Alice tell the company to find?

A couple of points:

–Assume the bidders have valuations randomly drawn from the uniform distribution on [500,1000]

–Suppose Alice holds an eBay style auction and she will sell the item for a price equal to the second highest valuation of the bidders*

(*this is a standard result in auction theory, though technically it's for one bid above the second highest valuation. An example: if bidders had valuations of $500, $600, and $700, the person who values the item at $700 would win the auction. The price he would pay in an eBay style auction with dollar bid increments is $601–just enough to outbid the person with the second highest valuation)

The puzzle is about two conflicting forces: Alice wants more bidders to bring her higher bids, but she faces a tradeoff in the cost of acquiring bidders.

Can you figure out the optimal number of bidders?

Answer to Puzzle 20: Finding the right number of bidders

Alice wants to maximize her expected auction profits. The equation for profits for *n* bidders is something like this:

Profit(n)= E(revenue n) – Cost(n)

The cost part is easy to figure out. Alice pays $10 per bidder, so her cost is 10*n*.

The harder part is figuring out the expected revenue for *n* bidders. What we want to know is the following. If we take *n* draws from a uniform distribution, what is the expected value of the second-highest draw?

This question is actually part of a larger topic in probability called order statistics. One can explicitly solve for the expected value of any distribution.

I will not go through the math here. But I will mention the order statistics for the uniform distribution are easy to visualize. What happens is that if you take n draws from the uniform distribution, the expected value of the n draws can be visualized as n points being evenly spaced on the interval.

Here is a picture to illustrate what I mean:

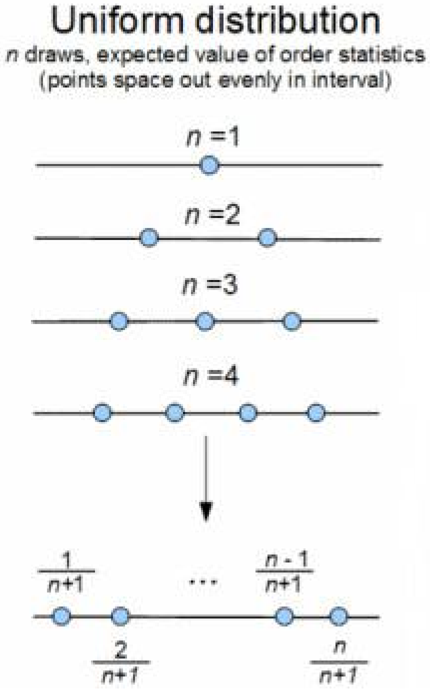

So the n points separate themselves along the interval. So you divide the interval into $n + 1$ segments, and the points will be at the fractions $1/(n + 1)$ along the way for the minimum, then $2/(n + 1)$ along the way for the second lowest point, etc., until the maximum draw which has an expected value of $n/(n + 1)$.

By this logic, the second highest draw is expected to be at $(n - 1)/(n + 1)$ along the way from 500 to 1000. This means the second highest valuation is expected to be:

$$500 + 500 * (n - 1)/(n + 1)$$

This is our formula for expected revenue. So we can substitute this expression back into the formula for profits:

$$\text{Profit}(n) = \text{E(revenue } n) - \text{Cost}(n)$$
$$\text{Profit}(n) = 500 + 500 * (n - 1)/(n + 1) - 10n$$

Now we need to solve for the profit maximizing point. I will skip the calculus and get to the point.

The profit maximizing point happens at n = 9 bidders, and Alice can expect $810 of profit.

The lesson is that more bidders is not always optimal: you capture much of the expected revenue from the first few bidders, and then the returns are diminishing (unless some bidder is a big outlier and you can extract money from him).

Extension: suppose Alice earned the highest valuation

As an extension, let's imagine that Alice somehow was able to extract the highest bidder to pay his entire valuation. This is not an assumption used in theory, but let's say it happens because of some irrational bidding war.

In that case, Alice would expect to earn slightly more revenue (the term $(n - 1)/(n + 1)$ becomes $n/(n+1)$), meaning her profit function is:

$$\text{Profit}(n) = 500 + 500 * n/(n + 1) - 10n$$

How will that change the number of bidders?

We can solve for the profit maximizing point and find that $n = 6$.

So Alice will only need to acquire 6 bidders, but she will earn nearly $870. This is 3 fewer bidders than above and she gets about $60 more.

This is, of course, exactly what we would expect: if Alice can extract more money from the bidders–the highest valuation instead of the second–she does not need as many bidders and in this case she can earn more profits.

This is common sense, but it’s useful to check the theory matches intuition.

More puzzles!

If you liked these puzzles, you will definitely like to read my blog Mind Your Decisions where I post a new math puzzle every Monday.

Table of Contents

Table of Contents

@Created by PDF to ePub