CHAPTER 9

The perfect rugby conversion

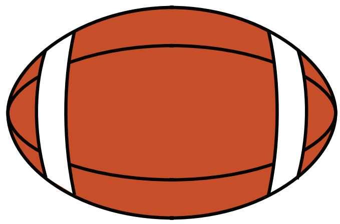

Yes, we know. Melbourne is an Aussie Rules town. But now we can all become passionate rugby league fans.[1] Australia is set to avenge their shock loss to arch rival New Zealand in last year's World Cup.

We've been discussing tactics with Cameron Smith, star kicker for the Melbourne Storm. To give the Kangaroos an extra edge against the Kiwis, we are figuring out a Maths-Masterly method of kicking goals.

As a reminder, a player scores a try by grounding the ball at a point in the *in-goal area* (that is, the end zone), beyond the opposition's goal line. The kicker then places the ball wherever he wishes along the red line pictured, through the grounding point and perpendicular to the goal line: we'll call this the conversion line. To score the extra points for the conversion, the ball must then be kicked between the goal posts and above the crossbar.

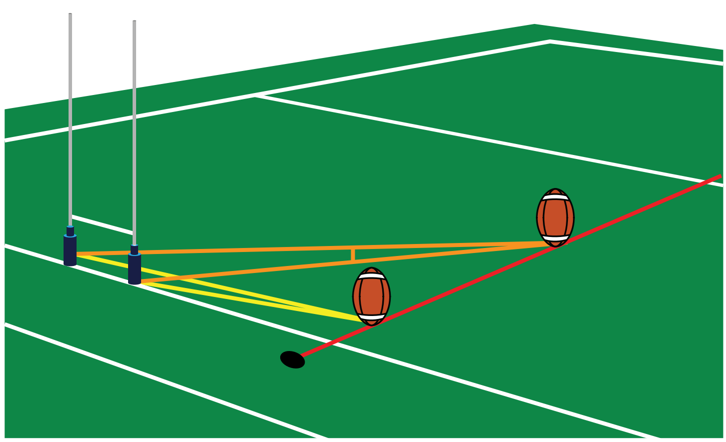

[1] This column was written in 2009. Rugby is a precursor to gridiron and is still popular in many countries, including parts of (but not all of) Australia. A *try* corresponds to a touchdown, and a *conversion* corresponds to an extra point.

But where should the kicker place the ball? If the ball is grounded between the posts, then the decision is pretty much a no-brainer. The kicker simply comes close to the goal line, subject to allowing enough distance for the ball to rise above the crossbar.

The method we describe deals with the trickier situation, when the try is scored to one side of the goals. In this scenario, the kicker should allow as much leeway as possible for skewing the kick. This amounts to making the angle subtended by the goalposts as large as possible.

But where is this spot? The angle is obviously zero at the goal line and then increases as we walk along the conversion line. Once we get to the other end of the field, the angle is again very small. The spot we're after is along the way, right where the angle begins to decrease.

A clever way to find the optimal spot is to draw circles passing through the bases of the two goal posts. One of these circles just touches the conversion line. That touching point turns out to be the optimal spot from which to kick.

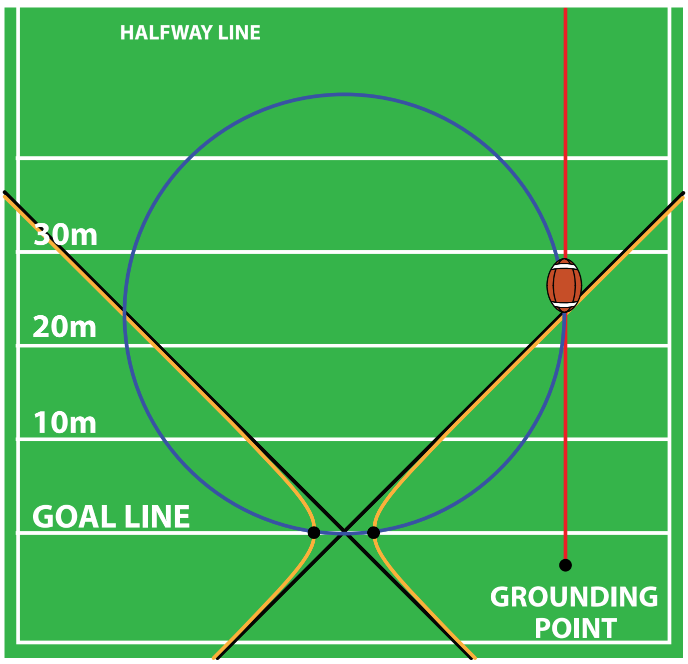

OK, it may be too much to expect the Kangaroos to come out drawing circles with their giant compasses. But there's another way.

If we take the optimal kicking spot from each conversion line, they combine to form the orange curve pictured. This curve should look familiar, since it is exactly a hyperbola. It is the graph $y = 1/x$, just rotated and scaled to fit on the field.

So, now we have rugby league players wandering out with their graphics calculators? Not necessarily. Notice that the hyperbola sits very close to its asymptotes, the crossing black lines.

These asymptotes hit the sidelines at about the 34 meter mark. So, the kicker can simply start there on the sideline and march straight towards the goals until he gets to the conversion line. That is then very close to his optimal spot. Simple!

There is one wrinkle to our calculations above. We have been drawing all our graphs down on the field, but the ball will hopefully be flying high over the crossbar. This changes the precise angles we need to investigate. Somewhat surprisingly, the optimal kicking spots are still along this same hyperbola.

We're finally done. And now what does Cameron Smith think of our brilliant mathematical plan? Alas, he simply prefers to kick from closer in, naively guided by factors such as deviation by the wind and curve of the ball in flight. And his intuition comes from having kicked a mere 341 career goals.[2]

We suspect Cameron knows exactly what he's doing and can comfortably beat the Kiwis without our help. It's back to the drawing board for the Maths Masters.

References: The following article, available at `www.qedcat.com/articles`, has all the gory details together with a comprehensive list of references:

B. Polster and M. Ross. *Mathematical Rugby*, Mathematical Gazette 94 (November 2010), issue 531, 450–464.

The article in which this mathematical way of locating the optimal conversion spot was highlighted for the first time is

Hughes, *A.*, *Conversion attempts in rugby football*, Mathematical Gazette 62 (1978), 292–293.

Puzzle to ponder

The method described above is based upon the following fact. Draw a circle passing through the two goal posts, pick a point on the circle, and draw lines to the bases of the two goalposts. Then, no matter the point, the angle formed by the two lines will be the same. Can you now fill in the details of the described method?

[2]Cameron Smith had kicked his 341 goals at the time of writing, in 2009. By mid-2017, this had risen to 1049 goals.

CHAPTER 12

Around the world in 80 (plus 129) days

In 2010, after 209 straight days at sea, Jessica Watson completed her trip around the world, alone and unassisted. It would be a terrific achievement for anyone, and 16-year-old Jessica is now the youngest person to have accomplished it, eclipsing the 1998–1999 effort of then 18-year-old Jesse Martin. *Except* the respected website Sail-World.com claimed that Jessica didn't actually sail around the world!

Was Sail-World correct? In attempting to untangle this dispute, we had to consider what it actually means to sail around the world. It is a surprisingly tricky question.

Almost everyone would agree that following a (relatively) small circle around the Antarctic would be cheating. A precisely mathematical way to rule this out is to demand that the circle be as large as possible—the 21,600 nautical miles for the route around the equator. (A nautical mile is about 1.8 kilometers, slightly longer than a regular mile). However, this would require Jessica to drag her yacht over large stretches of dry land, which seems a trifle unfair. There are other maximum-size circles one could try, but no such circle passes entirely through the oceans.

Anyway, it is clearly fussy and impractical to demand such a precise route. A more natural requirement, which is often suggested, is that the route finish where it begins, and that along the way the route must pass through two points on exactly opposite sides of the Earth. Such points are called antipodal, with the most famous antipodal pair being the North and South Poles (of course, a useless pair for Jessica).

A practical difficulty with this requirement is that the natural sailing routes lie low in the southern hemisphere, and all the antipodal points are located in the northern hemisphere. So, in order to pass through antipodal points, a low southern route must make a significant detour into the northern hemisphere. This is exactly the type of route Jesse Martin followed.

A map of Jessica Watson's route makes clear that, though her yacht did poke its nose across the equator, it never passed through antipodal points. So, there is a sense in which it is fair to say that Jessica did not sail around the world.

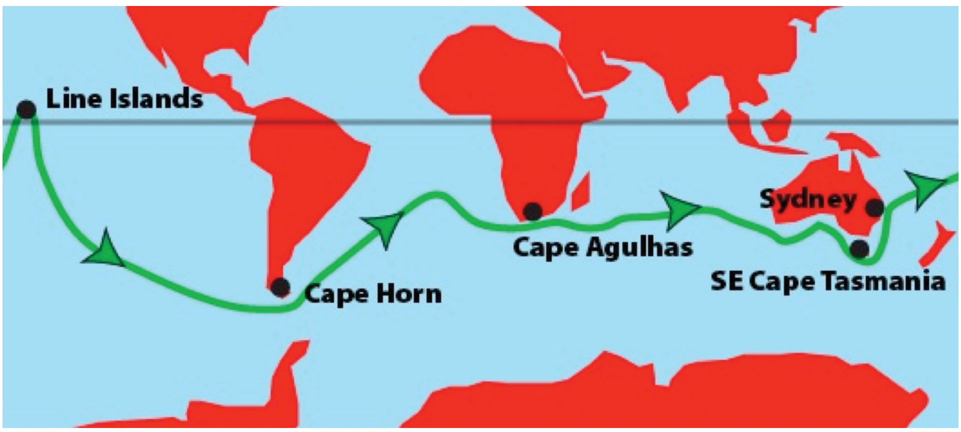

However, that is by no means the end of the story. In fact, the criticism from Sail-World has absolutely nothing to do with antipodal points.

Records for sailing "round the world" are governed by the World Speed Sailing Record Council, which has published a list of rules. To us, these rules seem quite arbitrary. For example, there is no requirement to pass through antipodal points, but the route is required to cross the equator (the reason for Jessica's northern detour). Why make people bother?

Nonetheless, the WSSRC can impose whatever rules they wish. The one rule of concern to Sail-World is *the shortest orthodromic track of the vessel must be at least 21,600 nautical miles in length.*

The 21,600 nautical miles is obviously there to ensure that the total distance travelled at least matches the length of an equatorial route. Fair enough. However, the *orthodromic track* is specifying how that distance must be calculated, and it is here that things get murky.

A path between two points is called orthodromic if it is the shortest possible path connecting the points. (Note that we're staying on the surface of the water here: submarines travelling along straight lines through the water are not permitted). But how does one use this to analyze Jessica's path?

To apply the distance rule, the reporters at Sail-World chose five land locations close to Jessica's route: Sydney, the Line Islands, Cape Horn, Cape Agulhas,

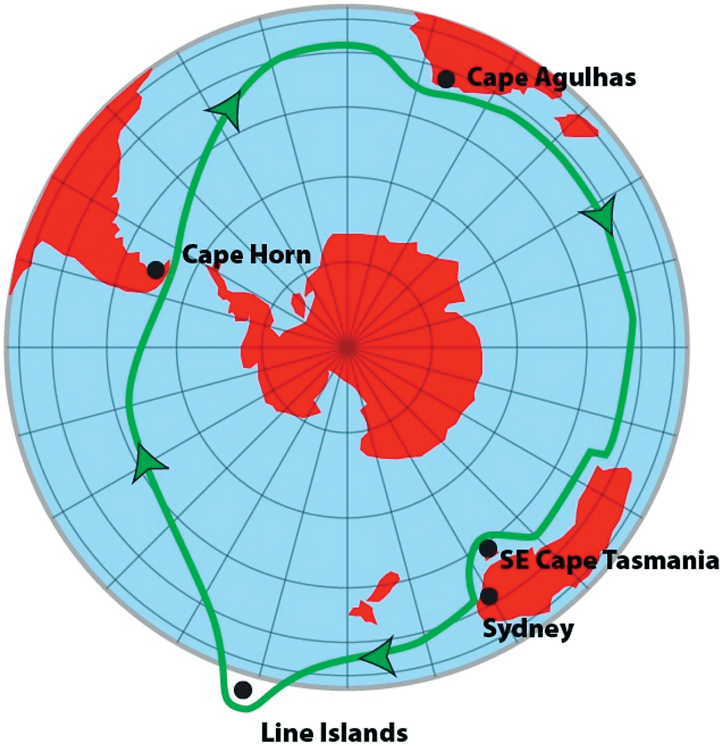

SE Cape Tasmania, and Sydney again to finish. The reporters then calculated the five orthodromic distances between adjacent points and summed these distances. They concluded that, even though Jessica's actual path exceeded 23,000 nautical miles in length, the orthodromic track of Jessica's route amounted to only 18,265 nautical miles, well short of the required distance.

We disagree. Even with the peculiar and confusing WSSRC rules at hand, we can make no sense of Sail-World's calculations. Firstly, the land locations seem arbitrary, since they only very roughly approximate Jessica's actual path. Secondly, we can see no reason to use only five locations, and using more would significantly increase the orthodromic track. In fact, using sufficient locations on Jessica's actual route would give an orthodromic track close to the 23,000 nautical miles that everybody agrees Jessica travelled. Finally, even if the WSSRC rules demand that the orthodromic calculation use land-based points (for heaven knows what reason), we believe that throwing in a few convenient islands near Jessica's route may raise her orthodromic track up to the required 21,600 nautical miles.

What if we're wrong, and Jessica's trip came up 3335 nautical miles short? As we read the WSSRC rules, Jessica could have made up the distance by sailing across Sydney Harbour a thousand or so times. Well, if that's what makes the difference between going "round the world" or not, who on Earth cares?

So, that's our conclusion. We believe the WSSRC rules are silly, and the complaints from Sail-World are sillier. To the extent that the requirements make any sense at all, we believe Jessica Watson has satisfied them. And, in the manner that most people seem to understand what going round the world means, we believe

Jessica has in fact done so. From the Maths Masters, and clearly from many others, congratulations!

Puzzle to ponder

We have suggested above that going around the world perhaps should require passing through a pair of antipodal points. Can you think of a silly route that goes through antipodal points but does not "really" go round the world? Can you think of a rule to exclude such silly routes?

CHAPTER 19

More forensicking

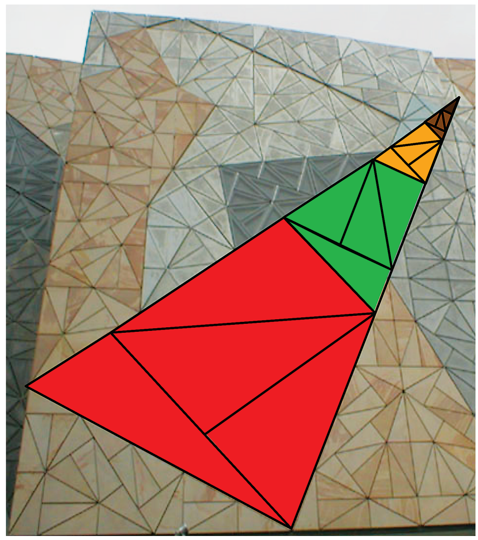

So, here we are, back in Federation Square. We've already taken a very close look at the square's triangular tiling, to see what the architects got right and (arguably) what they got wrong. However, we were inspired to write more by an unexpected encounter with Federation Square's doppelgänger. It appeared in an article by mathematician Roger Nelson,[1] author of the marvellous book *Proofs without Words* and its sequel.[2]

To explain this appearance, let's first recall the special property of Federation Square's façade. We begin with a big, judiciously chosen portion of the façade, as outlined in the picture above.

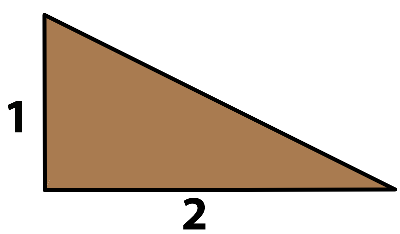

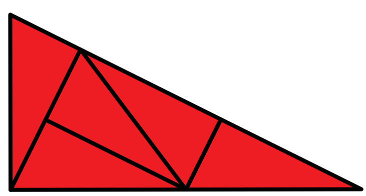

[1] *Proof without words: Right triangles and geometric series*, Mathematics Magazine, **79** (2006), 60.

[2] Mathematical Association of America, 1993 and 2001.

This triangular piece is just the right shape so that it can be replaced by five red sub-triangles of the same shape. Then those sub-triangles can be replaced by green sub-sub-triangles, and so on.

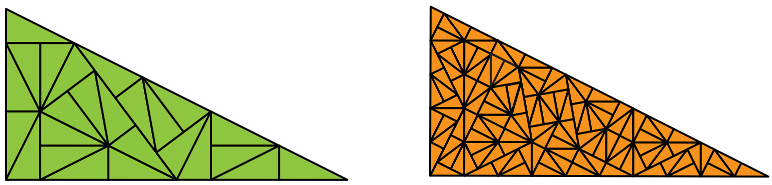

We can stop with the orange tiles pictured, giving us part of the Federation Square tiling. However this is when the doppelgänger appears, enticing us to keep going.

Let's begin again with the five red sub-triangles. However, at the next stage we only replace the rightmost red triangle.

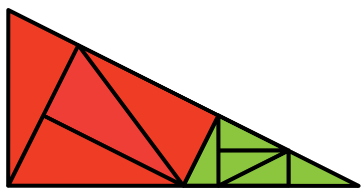

Similarly, at the next stage we only replace the rightmost green triangle.

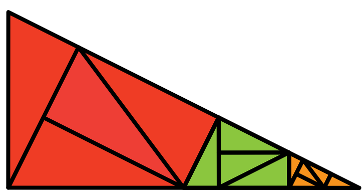

We do the same at the next stage, and the next, and... we just don't stop. There are infinitely many stages, where at each stage the rightmost minuscule triangle is cut into five micro-minuscule triangles.

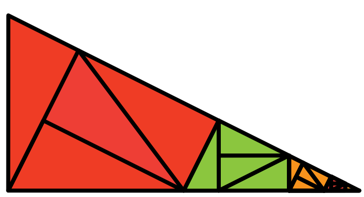

It's a cool picture but is there a point to it?

Recalling the mantra "one-half base times height", we see that the original brown triangle has area 1. It follows that each of the four red triangles have area $1/5$, the green ones have area $1/5 \times 1/5$, the orange ones have area $(1/5)^3$, and so on. Summing the areas of all the sub-triangles, our diagram shows at a glance that

$$\frac{4}{5} + \frac{4}{5^2} + \frac{4}{5^3} + \cdots = 1.$$

Or, after dividing both sides by 4,

$$\frac{1}{5}+\frac{1}{5^2}+\frac{1}{5^3}+\cdots=\frac{1}{4}.$$

That is a simple and beautiful summing of an infinite geometric series. Not surprisingly, other infinite series can be similarly summed, and Roger Nelson uses different right-angled triangles to calculate four such sums. For example, the diagram

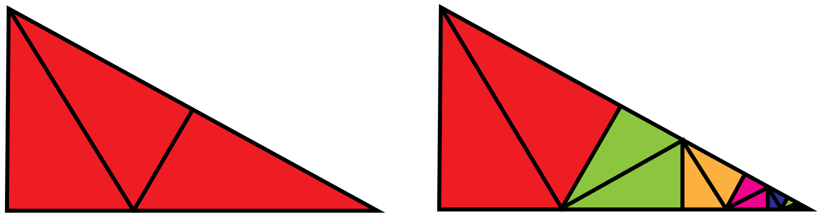

demonstrates that

$$\frac{1}{3}+\frac{1}{3^2}+\frac{1}{3^3}+\cdots=\frac{1}{2}.$$

You can probably guess the general formula:

$$\frac{1}{N}+\frac{1}{N^2}+\frac{1}{N^3}+\cdots=\frac{1}{N-1}.$$

This formula is true for any number $N>1$ (and any $N<-1$), and there is a pretty and well-known algebraic proof of the formula. However, we'll stick to geometry and to positive whole numbers.

We have taken care of $N=3$ and $N=5$, so can we find triangles to demonstrate the formula for other whole numbers N? It is not difficult to find right-angled triangles that work for $N=2$ and $N=4$, but not all numbers are so easy. Indeed, there is no triangle that works for $N=6$.

However, there is no reason to stick to triangles. The same argument would work for any figure that can be dissected into smaller figures (all of the same size) of exactly the same shape. Such figures were first studied by mathematician Solomon Golomb, who gave them the charmingly apt name of rep-tiles. They were then popularised by Martin Gardner in one of his excellent columns.

Some rep-tiles are not too difficult to find, but others are clever and beautiful:

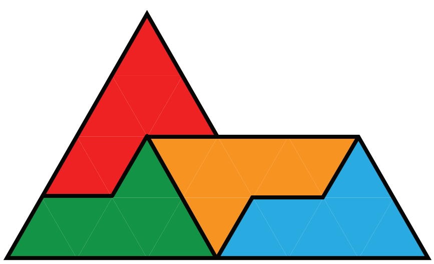

Of course, if we're willing to go down a dimension, then the simplest rep-tile is a plain old line segment.

Summing the lengths of the sub-segments, we obtain our first geometric sum (for $N = 5$), exactly as we did with Federation Square's triangles. Still, Federation Square does it with much more style.

Puzzles to ponder

A rep-tile that splits into N identical pieces is called rep-N. So, the tiles discussed above are rep-5, rep-3, and rep-4.

1. Find a triangle that is rep-2.
2. Find a triangle that is rep-4.
3. Given any positive whole number $N \geq 2$, find a rectangle that is rep-N.

CHAPTER 27

The super-rigging of poker machines

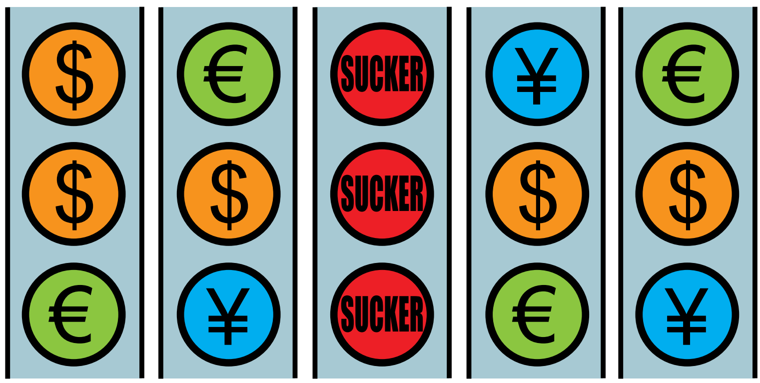

Are poker machines rigged?[1] Of course! Even the most optimistic gambler knows that the odds are stacked against them. But it is worse than that: poker machines are super-rigged.

Gambling is gambling. For those who engage, its unpredictability is part of the charm and the excitement. But for any gambling game, mathematics can still guide your expectations.

If a game is fair, then winning or losing comes down to luck, and in the long run you expect to get back close to what you gambled. Of course, no one expects pokies to be fair. In fact, of each \$100 gambled, you can only expect to get back about \$90.[2] It is in this sense that the pokies are rigged. This rigging is also more costly than it may appear, because it is likely that the same money will be gambled over and over, with a fraction lost each time.

But how can poker machines be super-rigged? The trick is in the psychology—to make it appear that the chances of winning are greater than they really are. Such scams have a proud tradition in carnival games. The comparable super-rigging of some poker machines has been documented by Melbourne barrister Tim Falkiner.

To illustrate super-rigging, consider the following simple game. Throw a pair of dice. If the sum is 10 or higher, then you win the Big Jackpot, with smaller prizes for other totals. This game can easily be rigged to give less than 100% money returned, simply by adjusting the sizes of the prizes.

Now consider the same game, but super-rigged with dodgy dice. Imagine we swap the 4 and the 6 on one die for the 3 and the 5 on the other: so, the sides of the dice are now 1-2-3-3-5-5 and 1-2-4-4-6-6.

[1] Australians refer to slot machines as poker machines, or pokies.

[2] In Australia a minimum long-run return of around 90% is guaranteed by law.

The dodgy dice together have the same numbers as normal dice. This implies that the average sum of a roll of the dice has not changed, giving the psychological impression that prize-winning totals are just as likely.

But actually, the chance of winning the Big Jackpot has decreased dramatically. With normal dice there are six of the 36 possible rolls that will sum to 10 or higher; the rolls are indicated in the left diagram below.

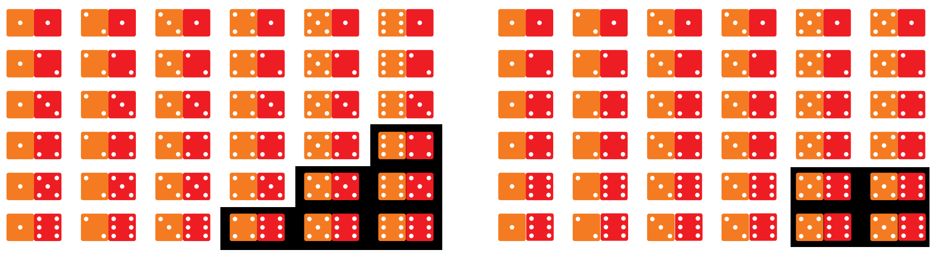

With the dodgy dice, there are now only four such rolls, as indicated on the right. So, the chances of winning the Big Jackpot have gone down from 6/36 to 4/36—a 33% reduction. And, cunningly, the chance of "just missing" by rolling a 9 has doubled, from 4/36 to 8/36.

It is exactly this type of super-rigging that is programmed into the pokies. Think of the five poker machine wheels as 30-sided dice. Some wheels are starved of Jackpot symbols, which is then disguised by loading a few more Jackpot symbols on other reels. The consequence is that winning the Big Jackpot is much less likely than it appears. And, the chances of "just missing"—encouraging another go—is much more likely.

The use of super-rigged poker machines is incredibly sneaky. Is it also illegal? Tim Falkiner has argued that if such machines are "deceptive", then they could well be banned under the Trade Practises Act.[3] However, the Australian Competition and Consumer Commission is apparently unconvinced that such machines are deceptive. We feel compelled to ask: what would be?[4]

Puzzle to ponder

It is quite well known that on any given roll of two dice, a total of 7 is most likely. Now imagine you are playing the following game: A pair of dice is thrown over and over. You want a total of 7 to appear twice. Your opponent is waiting until a total of 6 and a total of 8 have appeared. Would you bet that you can get there first?

[3] An Australian consumer protection law, which in many instances has turned out to be very powerful. Of course, America has no comparable history of consumers being fleeced by dodgy companies, and so Americans have no need for such a law.

[4] In late 2016 a pokies player brought a lawsuit against a casino and poker machine manufacturer, arguing this very point. At the time of writing, the case has not been settled.

CHAPTER 40

Sex, lies, and mathematics

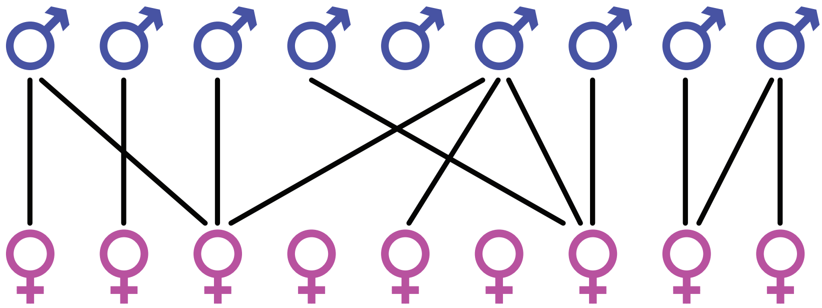

Statistical surveys regularly report that, on average, heterosexual men have more sexual partners than heterosexual women. Consider, for example the survey included in the "Sex in Australia" series, published in 2003 in the *Australian and New Zealand Journal of Public Health*. This was a very large survey, which included 9469 men and 9340 women who identified themselves as heterosexual.

The responses of the heterosexual men indicated an average of 3.9 sexual partners in the previous five years. By comparison, the responses of the heterosexual women indicated an average of only 1.9 partners, less than half that of the men. This difference is amazing. And, for Australia as a whole, it is mathematically impossible.

Let's place all the men and women in two rows and then draw lines between all the "friendly" pairs as shown above. To calculate the average number of partners for the men, we sum the partners for all the men—which amounts to counting the number of connecting lines—and then divide by the number of men. Similarly, we can calculate the average number of partners for each woman.

But of course the number of connecting lines is exactly the same for both calculations. And, as the survey implies, there are about the same number of heterosexual men and heterosexual women in Australia. So the averages must be very close to identical!

The authors of *Sex in Australia* remark upon similar differences in surveys from around the world. A suggested explanation is that women are more accurate in the way they count partners, with (surprise, surprise!) men overestimating their partners. A related factor is the prevalence of sexual double standards: people can simply be too embarrassed (or too macho) to tell the truth.

Of course, such huge discrepancies muddy the true figures, making such reports of limited use. And the message we get is likely to be even less useful, since the media tend only to report these impossible averages.

Math has alerted us to the problem, but how do we solve it? How can we encourage people to tell the truth in such sensitive surveys? Maths may help with a solution as well.

Suppose we are surveying a group, and we want to find out how many within the group were friendly with someone during the previous week. OK, so you ask the question. But, you also tell everybody to make their response dependent on tossing two coins.

The person is instructed to lie if two heads come up and otherwise to tell the truth. Then everyone can tell the embarrassing truth without anyone else being certain whether they are doing so. Similarly, there is now little point in bragging about all your friends. But do these coin-manipulated responses tell us anything? Yes!

Suppose there are 10,000 people honestly using the coin method to respond to our question, with 6000 replying “Yes” and the rest replying “No”. Then some simple mathematics shows that there were about 7000 people who were friendly last week. We’ll leave you to figure out why.

Reference: As far as we know, the simple argument presented in this chapter is due to David Gale, who used to be a professor of mathematics at the University of California at Berkeley. He is also one of the mathematicians responsible for the *stable marriage theorem*, discussed in the next chapter.

Puzzle to ponder

Explain our estimate of 7,000 friends.

CHAPTER 44

The Maths Masters' Tour de Victoria

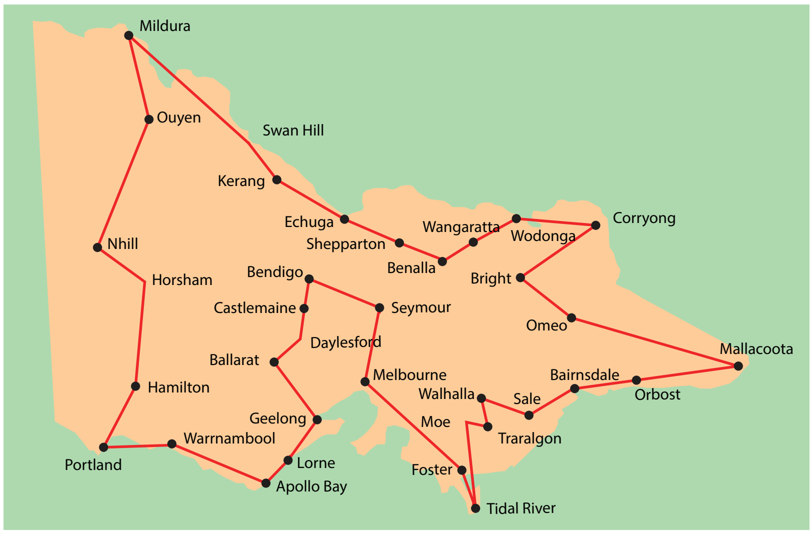

Surprising as it might seem, mathematical speakers are sometimes in very high demand. In anticipation of things going really crazy, we have planned the Ultimate Maths Masters Tour. Our route will take in 34 of Victoria's friendliest and prettiest towns.

The plan is to start in Melbourne and to fly the Mathscopter in straight lines from town to town. Our route will consist of one big loop, finally ending back in Melbourne. There are many loopy routes to choose from, but being Maths Masters we can only be really happy with the shortest possible one.

We are reasonably sure that the shortest tour of Victoria is that indicated in the diagram, coming to a total of 2172 kilometers. Why only reasonably sure? After all, isn't it simply a matter of taking all possible loops and measuring which one is the shortest? Yes. And No.

Starting in Melbourne, there are 34 possible towns for our first stop. There are then 33 possibilities for the second town, then 32, and so on. This means that we have $34 \times 33 \times 32 \times \cdots \times 3 \times 2 \times 1$ possible loops to consider. Oh, but we do get to divide by 2: this monster multiplication counts each loop twice, once for each possible direction of travel. That leaves us with a mere

$$147,616,399,519,802,070,423,809,304,821,760,000,000$$

loops from which to choose.

That's a lot of loops. Imagine we had a billion supercomputers working away, each analyzing a trillion loops per second. Then we'd have all the loops analyzed in about five billion years: probably just in time to see Victoria and the rest of the Earth plummet into the sun, making the whole calculation somewhat redundant.

So, how did we come up with our proposed route? The argument comes in two parts, beginning with a very clever idea. To illustrate, imagine starting with any loopy tour, and choose any two unconnected segments of that tour. They may be anywhere, but to begin with, let's suppose the two segments actually cross.

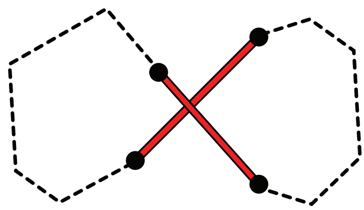

We now consider replacing those two segments with two different segments connecting in pairs the same four towns: there will always be such a choice of segments that also keeps the whole tour as one big loop. And, since the original segments crossed, the new segments will be uncrossed and will definitely shorten the tour.

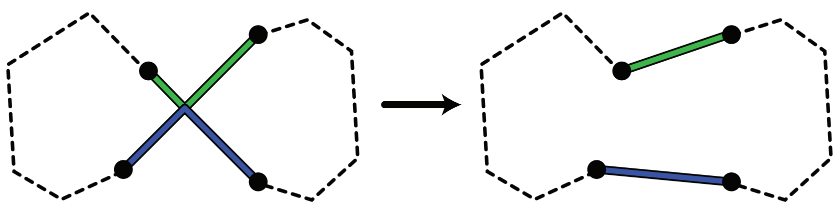

Even if the original segments do not cross, the competing segments may shorten the tour. It's simply a matter of trying and comparing.

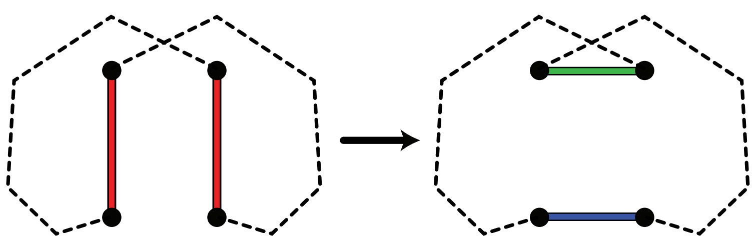

With this simple idea, we now have an easy and relatively quick procedure for coming up with a candidate for the shortest tour. We first start with any loopy route, and we consider all possible pairs of unconnected segments: for our Victorian tour, a total of 561 pairs.

For each pair we then consider whether the suggested interchange shortens our tour. We choose whichever interchange shortens it the most, creating a new loopy tour. Then we try to shorten again, and again, until no interchange shortens further. This final route is our candidate for shortest route overall.

This procedure does not always detect the truly shortest route. However, in practise the approach often gives strikingly good results, very quickly providing a route which is very close to the shortest.

What we have been discussing here is known as the *travelling salesman problem*. The difficulty of this problem, of guaranteeing to have found a truly shortest route, is infamous. No one knows a watertight procedure that does not take an astronomically long time. So, much work is devoted to "reasonably sure" procedures of the type we have described.

Very annoying! Mathematicians hate to say "near enough is good enough". But, when the alternative is to wait billions of years for the answer, sometimes even mathematicians are willing to bite their tongues.

Puzzle to ponder

Let's say there are 64 cities located at the centres of the squares of a chessboard. Find a shortest round trip that includes all cities.

CHAPTER 48

Maths MasterChefs

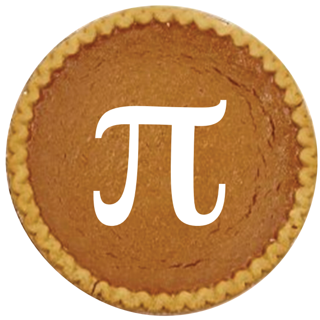

Cooking shows are all the rage. It seems that every night television chefs compete for culinary honors, and that has given us an idea for another reality show: *Maths MasterChef.*

We want to give our show a trial run. So, here are some of our tastiest mathematical dishes, and you can judge how you think they'll fare. Perhaps they could even add some nutrition to current school offerings.

Entrée: Nourishing Nines Niçoise

For this recipe, take the elusive number $0.999\ldots$, and prepare the equation $M = 0.999\ldots$. This will require infinitely many 9's, so try to buy them on special. Now, gently multiply both sides of the prepared equation by 10. If you were careful, you should find that $10M = 9.999\ldots$. Firmly subtracting the first equation from the second, the result is delectably simple:

$$\begin{array}{rr} & 10\,M = 9.999\ldots \\ - & M = 9.999\ldots \\ \hline & 9\,M = 9.000\ldots \end{array}$$

Now, quickly divide by 9 and you will discover that M is also equal to 1. So, $0.999\ldots = 1$. Very tasty!

Main Course: Traditional Triangle Casserole

To prepare this Greek classic, you will need four red right-angled triangles and a blue square pot.

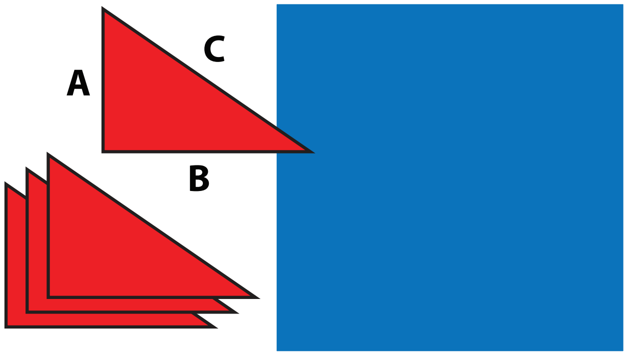

Presentation is everything, so make sure the triangles are identical and are just the right size to fit neatly around the edge of the pot.

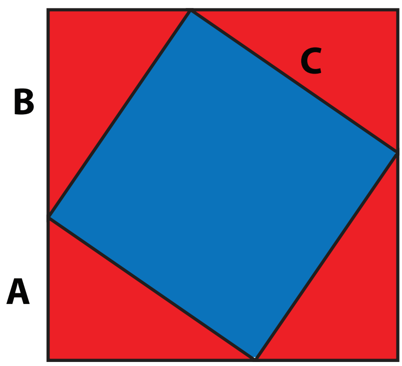

Then the uncovered area of the pot is a square of side length C. That gives an enticing blue region of area C^2. Now, stir the triangles around a bit.

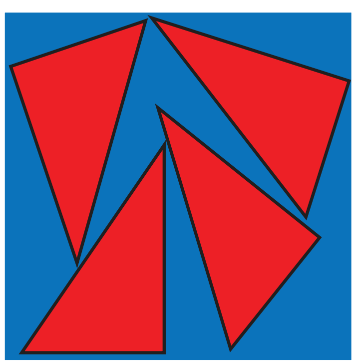

What is our uncovered blue area now? It is C^2, the same as before. Stir again, very carefully...

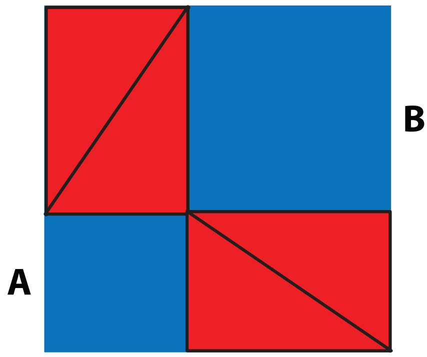

The uncovered blue region still has area C^2. However, the blue region now consists of the two little squares. So, $A^2 + B^2 = C^2$. A wonderfully satisfying dish, a gift of that legendary mathematical gourmet, Pythagoras.

Dessert: π-can Pie

Take an appetizing circle of radius R, and slice it into an even number of identical wedges. The circumference is a delicious $2\pi R$. Now, rearrange the wedges into a parallelogram.

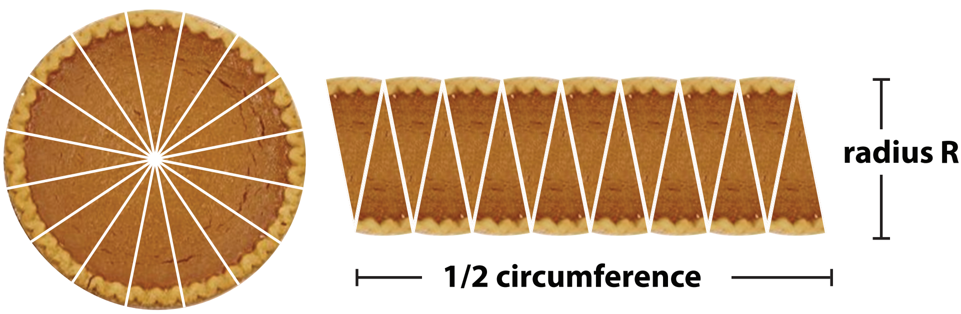

Then the area of the circle is 1/2 of its circumference times its radius R. Very sweet! Combine this with a previously prepared circumference of $2\pi R$, and you can serve up the area of the circle to be πR^2. What a smart dessert to finish a mathematical meal! (A tip for perfectionist chefs: the more wedges and the thinner they are, the more satisfying this dish will be.)

So judges, what do you think? Will such recipes rocket *Maths MasterChef* to the top of the TV ratings?

Puzzle to ponder

And, here is something you can try using leftovers from preparing the Greek dish above. Take two red triangles and place them in a blue rectangular pot. Stir the triangles until you can see at a glance that the blue parallelogram has the same area as a rectangle with the same base and height.

CHAPTER 57

The statistical problem of greedy pigs

Your Maths Masters love games. We're always ready for a sophisticated game of chess, or Go, or Snakes and Ladders. Yep, your Maths Masters are not too old to enjoy games of pure chance, letting the gods of the dice decide our fate. And, we very much enjoy games that require both luck and skill, which makes us big fans of Greedy Pig.

For those unfamiliar with it, Greedy Pig is a simple children's game, played over five rounds. In a given round a die is rolled repeatedly, with each player's total increasing by each number rolled, *unless* the "killer number" 2 appears. To avoid being killed, a player can choose to stop before any roll, locking in their total for that round; the round ends when all players have locked in or when a killer 2 appears, at which point all the greedy players who haven't locked in their scores lose all their points for that round. The winner is the person with the highest score over the five rounds. (Needless to say, there exist many minor variations of the above rules.)

Greedy Pig is a fun game with easy rules. It's also an excellent introduction to the intuition and ideas of probability, making it very attractive as a classroom exercise. So, how does one go about investigating Greedy Pig in a mathematical manner? How does one become Greedy Pig World Champion?

A detailed lesson plan for Greedy Pig is available from the very popular maths300 education website, and a shorter plan in the same style can be found at the NSW Government's Digital Education Revolution (DER) website. As maths300 describes it, the goal is to try to determine the "best strategy" for Greedy Pig. At which point maths300's lesson has already gone somewhat off the rails.

We're not trying to be nitpicky. (If we were, we'd point out to the maths300 writers that the singular of "dice" is "die".) However questions about probability can be very subtle, which means if we're going to be searching for a "best strategy", we really must know what that strategy is supposed to achieve. What is the actual goal? Unfortunately, maths300 never says.

In fact, there is no general "best strategy" for playing the *game* of Greedy Pig. Math can help, but winning Greedy Pig will depend more upon psychology, upon evaluating the strategies of your opponents. For example, if you're playing against a number of greedy players, then playing "sensibly" may result in you beating most of them, but it is very unlikely to make you the overall winner: chances are, a few lucky opponents will choose just the right rolls on which to be greedy. (Exactly the same difficulty arises in trying to win football tipping competitions.)

OK, so if the maths300 lesson isn't about searching for a best strategy to win the game of Greedy Pig, what then is the goal? What more straightforward mathematical goal might there be?

The point (which is implicit in maths300's lesson) is to not consider Greedy Pig as a game against opponents, but rather as a solo game. The individual player is then looking for a strategy to give a good chance of getting a high score. That is, the player wants to try to make her score *on average* as high as possible. This is now a precise mathematical goal, and we can get to work searching for a best mathematical strategy to achieve that goal.

As a lesson, Greedy Pig is intended to be very open and exploratory. Students are encouraged to come up with their own strategies, games are played, and outcomes are compared. Then, hopefully, there is extended discussion of why some strategies may be better than others. Why, for example, is "stop after three rolls" not a great strategy?

All this seems natural and good. The lesson is fun and it may well result in some students gaining some genuine understanding of the way probability works, without resorting to formulas or heavy calculation. If students simply learn that dice don't have memories, that every roll provides the same 1/6 chance of killing everyone and no matter how long since the last 2, that would be great. It would put the students well ahead of the thousands of foolish victims blowing their savings on the pokies.[1] Unfortunately, maths300 (and DER) goes seriously off the rails.

Exploration and experimentation is all well and good, but so are actual answers. So, what is the actual best strategy for averaging a high score on Greedy Pig? How can we determine that strategy?

After all the games, maths300 concludes its lesson with a computer. They decide to "[use] a computer to find the optimal strategy". The various strategies are programmed into the computer, some simulations are run, and the winning strategy is declared.

Oh dear.

Let's put maths300's computer aside for a moment and get back to actually thinking about the problem. Suppose you happen to have scored 15 points on a round, and you want to decide whether it's a good strategy to continue. Well, there's a 1/6 chance that on the next roll your score will go up by 1, and a 1/6 chance your score will go up by 3, and so on; *and*, there's a 1/6 chance you'll be hit by the killer 2, meaning your score will go down by 15. That means that on average your total on the next roll will change by

[1]See Chapter 27.

$$\frac{1}{6}(1+3+4+5+6-15)$$

Since the above sum is positive, the conclusion is that if you've only obtained 15 points, then (on average) it's worth risking another roll. And what if, more generally, you've already scored a total of T points? Then, on average, the next roll will change your total by

$$\frac{1}{6}(1+3+4+5+6-T)$$

All that matters is whether this sum is positive or negative. The conclusion is that you should keep going until your total on the round is 19 or above, and then stop.

Seriously, how hard was that? How can maths300 provide 15 pages of lesson plan without a single mention of this simple and conclusive arithmetic approach?

There seem to be two underlying reasons for maths300's approach, neither all that convincing. The first reason involves the intended audience: the maths300 lesson is supposedly designed for classes from year 3 to year 12, and "attempted explanations involving averaging and long-run frequencies call upon sophisticated concepts that the student may not yet possess".

Fair enough, up to a point. One would of course wish to be gentle with year 3 students, and the formulas above would be too much. However at some level, *well* below year 12, the simple arithmetic just has to be presented. Indeed, the T formula could be a very natural and impressive introduction to the power of algebra.

Moreover, though the notion of long-term average may be (somewhat) sophisticated, it is also intrinsic to an understanding of the problem, of what the "best strategy" is supposed to achieve. Shying away too much from such ideas makes the whole exercise meaningless.

Bizarrely, maths300 is even reluctant to mention any notion of probability, that the killer 2 has a $1/6$ chance of occurring on each roll. Supposedly such an approach "can be mystifying for a student who self-evidently has seen that the 2 sometimes occurs twice in a row and at other times might not occur for 20 rolls". Well, yes, it can be mystifying. It's also exactly the kind of mystification that a classroom lesson on probability and randomness should be confronting head-on.

The second reason underlying maths300's approach is even more concerning: a fundamental hostility to theory and the very notion of proof. The maths300 writers contrast "empirical versus theoretical" methods of learning and teaching, and clearly have no respect for the latter. For maths300 it seems that mathematics, whether in year 3 or year 12, is nothing more than an experimental science. The notion that a clear statement of probability or a simple formula for an average may be valuable or may solve the problem in a clean and conclusive manner, seems to carry no weight. (maths300 does offer a "theoretical approach" in an afterthought "Answers" document; however, the probability calculations there include none of the simple approach we've outlined above and pretty much answer nothing of interest. It also raises the further question of whether the writers know or care about the difference between a probability and an average.)

Your Maths Masters definitely appreciate the role of computer simulation in teaching Greedy Pig. The repeated rolling of a die can only provide intuition up to a point; after that, to see what "long-term" is really like, one has no choice but to trust the computer and see what happens with zillions of simulated rolls. Moreover, a game needn't be much more complicated than Greedy Pig so that computer simulation is the *only* way to get a reasonable understanding of the probabilities involved.

However Greedy Pig *is* a simple game, and it *does* have a simple, and simple to prove, solution. It is absurd to not present this solution to any student, at any level, who has some chance of understanding it.

Your Maths Masters have heard good things about maths300, that many of their lessons are well worthwhile. Maybe, maybe not. However, in the case of Greedy Pig, it appears that a dubious educational ideology has won out over simplicity and mathematical common sense.

Postscript: This column was pretty much the final straw for us, convincing us that Australian mathematics education was beyond repair. The column received (by our little column's standards) a large amount of heated criticism, none of it worth a dime. None of the many maths300 fans could face up to the fundamental absurdity of the way the game had been presented. Then, in 2016, the writers of maths300 contacted us, indicating they had altered their Greedy Pig lesson in response to our criticisms. However the lesson remained (and last we checked was still) a constructivist swampland, just with the formula solution thrown somewhere in the middle.

Puzzle to ponder

Suppose the rules of Greedy Pig are changed so that on your second roll you get twice the total on the die (except for the killer 2), on the third roll you get four times the total, on the fourth roll 8 times the total, then 16 times, and so on. When should you stop?

CHAPTER 62

Shooting logarithmic fish in a barrel

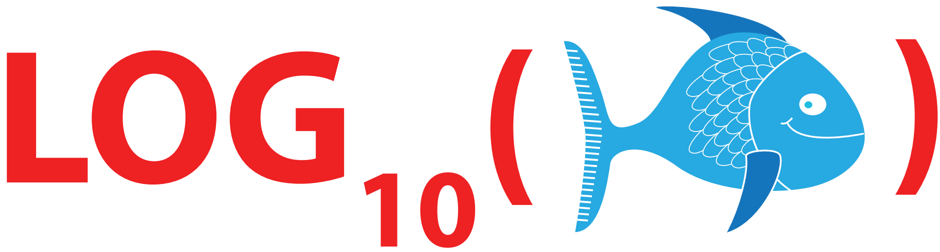

It should be clear by now that our columns vary considerably in style and content: on occasion we have worked hard to try to explain some pretty mathematics, at other times we've been happy to sit back and take easy potshots at mathematical nonsense. Both types of column have been fun and rewarding, at least for us (and we hope for others).

One potshot column that we particularly enjoyed putting together concerned the SAFE index,[1] a naive approach to measuring the vulnerability of endangered species. Though undoubtedly well intentioned, the SAFE index is an archetypal example of mathematical fancy dress camouflaging the thinness of the underlying ideas.

An unexpected benefit of writing the SAFE index column was our meeting Associate Professor Michael McCarthy. Michael is an ecologist at The University of Melbourne, working on conservation biology and ecological modelling, and he clearly grows impatient with some of the fluff that appears in his field. Michael and his colleagues (and others) had submitted their own, detailed critique of the SAFE index. Michael has also performed yeoman's work analyzing the (conservative) Victorian Liberal Party's ground-breaking study into the effect of cows trampling on fragile alpine vegetation.[2]

Michael also alerted us to a new venture, the Ocean Health Index (OHI). True to its name, the OHI is intended to be a measure of the health of the Earth's ocean ecosystem. It was launched in 2012 with an article in the prestigious journal *Nature* and is supported by a large international team of respected marine scientists.

So, the OHI has to be a good thing, right? Perhaps. However, Michael and a number of his colleagues have some pretty large nits to pick. (*Nature* declined to publish their letter in response to the original article, though an independent critique and authors' reply subsequently appeared in the journal.)

[1] See the previous chapter.

[2] Reactionary idiot cowboys permitting grazing in fragile terrain, with the paper-thin pretense of "studying" the effects. Australia's version of Japanese "research" on whales.

The OHI attempts to measure the status of ten public goals for a healthy ocean. Included are specifically physical and biological goals, such as water cleanliness and biodiversity, as well as human-centred goals, such as tourism. Each of the ten goals are scored out of 100, and then the scores are averaged to give the overall OHI. The Index is measured for each country as well as for the world overall. For example, Australia's OHI in 2013 was 70 (43rd in the world), compared to the Global Index Score of 65.

What's wrong with that? Plenty, and some of it obvious. Michael is diplomatic, suggesting that the meaning of ocean health is "unclear". Your less diplomatic Maths Masters will be more blunt: the Ocean Health Index is pointless.

To begin, one might conceivably attempt to distill the state of "water cleanliness" or "tourism" to a single number, but then what? If water cleanliness receives a score of 90 and tourism a 10, what can the average of 50 possibly indicate? What if the numbers were reversed?

There is clearly little meaning in the average of ten such numbers. Moreover, the extensive interrelatedness of the goals being scored muddies the little meaning that might exist. It follows that the Ocean Health Index is not a true measure of anything; it is just faking simplicity.

Furthermore, as Michael and others have noted, the ten individual scores are in themselves problematic, and intrinsically so. The OHI has the laudable purpose of measuring ocean health with respect to human interaction. However, the computational effect is that each score is an intermingling of physical and biological measures of the current ocean state with human pressures upon that state. The result is that even the individual scores are scores of nothing in particular, neither fish nor fowl.

It gets sillier. Michael has alerted us to some truly bizarre arithmetic lurking in the Index.

One of the ten goals in the OHI is food provision, being the "the amount of seafood captured or raised in a sustainable way". This is calculated as the weighted average of x_{FIS}, a score of "fisheries" health, and x_{MAR}, a "mariculture" (ocean fishing) score.

The *Nature* article does not indicate how x_{FIS} and x_{MAR} are calculated, however some details are provided in the 2012 supplementary materials. We're simplifying, but x_{FIS} is essentially a ratio, a comparison of the number of fish harvested in a year to the maximum sustainable harvest. To calculate x_{MAR}, one first calculates the number of fish caught divided by the area of the region being fished. (The detailed computation involves the consideration of different fish species and the incorporation of estimates of sustainability.) Then, one takes a logarithm, giving in effect the following formula:

$$x_{MAR} = \log_{10}\left(\frac{\text{Fish}}{\text{Area}} + 1\right)$$

Why the +1? God only knows. (In 2013 the +1 was eliminated from the formula.) Why the logarithm? Again, God only knows. And then, what to make of the overall food provision score: How does one average x_{FIS}, a unitless ratio, with x_{MAR} the logarithm of (1 more than) a fish density? We doubt that even God knows the answer to that one.

Mathematical modelling is difficult and subtle; one should avoid declaring hard and fast rules, and one should be wary of being too, or too quickly, critical. However, it is a fundamental rule to not sum quantities of different physical types. For example, if a car travels 200 meters in 7 seconds, we don't then create a Frankenstein quantity of 207 met-secs. Furthermore, though logarithmic scales can be enlightening, we have never seen logarithms employed in the manner attempted in the Ocean Health Index.

We are sure, despite its glaring flaws, that the Ocean Health Index is well intentioned and incorporates good and important research. It may be salvageable. However, we cannot understand why the creators of the Ocean Health Index chose to include such clunky mathematics in their model. Nor can we understand why *Nature* chose to accept and publish the article in that form.

Whatever the explanation, for us the expression "appeared in *Nature*" no longer has such an authoritative ring.

Puzzle to ponder

We've got a new idea for an index which measures what kind of "shoe person" you are. Our shoe index is the logarithm (base 10) of the average of your shoe size and the number of pairs of shoes you own. One Maths Master has shoe size 8 and owns two pairs of sandals. This gives your Maths Master an average of 5 and a shoe index of approximately 0.7. His wife has a shoe index of 2.3. How many pairs of shoes does she own?