The Maths of Life and Death
Dr Kit Yates

4 May 2021

Hi, my name is Kit Yates. I'm a Senior Lecturer in mathematical biology at the University of Bath and it is my pleasure to be giving this Gresham lecture this evening on The Maths of Life and Death.

I'm going to break the talk down into three main parts: first part is going to be looking at the places where mathematics says impact in the medical arena and places where a bit of a better knowledge of mathematics can really help us out; Secondly I'm going to look at places where maths appears in newspaper headlines on TV on the radio so maths and media; and third young women look at maths in the criminal justice system and try and understand places where maths can crop up in the context of crime.

Maths of Medicine

Firstly, to talk about medicine. I'm going to start off with a little problem which you can have a go at home if you'd like to and it's going to be a problem about screening. I'm going to talk to you about screening programmes in the United Kingdom specifically and I'm going to talk specifically about breast cancer although this is similar for other areas. I'm going to give you some statistics here's one of them which is about the probability that a woman over 50, who regularly gets invited to the screening programme in the UK - we invite women between 50 and about 72 or 73 for regular screens every three years - the prevalence of breast cancer is as low as 0.4% so that means for1000 women, randomly selected from the population, over the age of 50, four of them will have undiagnosed breast cancer. Another statistic that I'm going to give you is about the screening test itself. If a woman has breast cancer, the probability that she tests positive is 90% - so nine times out of 10 if you have breast cancer and you go for a screen you will be told to come to follow up tests one time out of 10 unfortunately your breast cancer will be missed. Another statistic is about the number of times that women who don't have breast cancer will get the correct results. Actually, a woman that doesn't have breast cancer goes through screen, the probability that she incorrectly tests positive - so gets a false positive - is 10%. But again 9 times out of 10 she will be correctly told that she doesn't have breast cancer.

Here are three statistics about the background prevalence of the disease in the population, which is quite low, and the probability that if a woman has breast cancer that she tests positive and if she doesn't [have breast cancer] the probability that she tests positive. I'm going to ask you a question that was asked to a number of German doctors few years ago - and this is the question that will see if you can do any better than they did - so “Which of the following best characterises the probability that a patient with a positive monogram actually has breast cancer?” Which of the following best characterises the probability - so I’m going to give you some numbers which will I'll put down as percentages, they are going to characterise probabilities – that a patient with a positive mammogram - someone get centre letter saying “you need to come from further follow up tests” - actually has the disease? I’m going to show you 5 numbers and you can choose between them:

Is it A) 90.0%? Is it B) 81.0%? Is it C) 49.1%? Is it D) 3.5%? or is it E) 0.4%?

Just to recap, the probability that a woman over 50 has breast cancer - undiagnosed breast cancer - in the UK is about 0.4%. If a woman has breast cancer and she goes through screen, the probability that she tests positive is 90%, so of nine times out of 10 she gets the correct result and one time I out of ten that breast cancer was missed. If a woman doesn't have breast cancer, the probability that she incorrectly tests positive is 10%, so again 9 times out of 10 shall be correctly told she doesn't have breast cancer and one time out of ten be given the incorrect false positive diagnosis. What's the probability that if someone gets a positive mammogram that they actually have breast cancer? I'll give you a second to think about this. If you're watching this afterwards then you can just pause the video, but I'll just pause for 10 seconds to allow people to think.

[Ten second pause]
Okay so the correct answer - and I think this is a surprise to many people when they first see this problem, it was certainly a surprise to me when I hadn't thought about this rigorously and mathematically – is D) it's just three and a half percent. That seems really alarming. Is it possibly true that of all the letters that get sent out are only three and a half percent of them are actually correct positive diagnoses? Where does this figure come from? Well, I'll tell you, but before we go into the detailed maths, I'll tell you that of these German doctors - I think they we're all actually gynaecologists - fewer than 20% of them got the answer to this question right, which suggests that, actually, they would have been doing better purely guessing at random in this question. So, it is a difficult problem. If you got it wrong there's absolutely no shame that - I didn't give you much time to work it out mathematically, but let's think about how we can actually get a handle on this figure.

Rather than playing around with percentages and decimals, I want to think about 10,000 representative women who get tested in this over 50 population. Of those 10,000 women I said that 0. 4% of them will have breast cancer - undiagnosed breast cancer, so that's 40 cases in this 10,000 population. The remainder, the vast majority - 99.6% - that's 9960 of these 10,000 women, so almost all of them will not have breast cancer. But all of them still get asked to go for a screen, so they all get tested, and they get a positive or negative result back. Of the women that have breast cancer, 90% of them will be correctly identified as having breast cancer, so that's 36 true positive results - 90% of 40 is 36. Unfortunately, 10% will be given a false negative diagnosis. That's four women giving a false negative diagnosis. Of the women that don't have breast cancer, again the vast majority will be correctly told they don't have breast cancer. That's 90%, that's 8964 true negative result. But because the size of this second population is so large, even a 10% false positive rate - 10% of 9960 is 996 - that's a huge number – that's nearly 1000 false positives. You can see that that massively outweighs the number of true positive diagnosis. Actually, the number true positives massively dwarfed by the number of false positives. If I work out the number of true positives as a proportion of the total number of positives - the proportion of correct positives is just 36 / (36 + 996) the total number of positive results - total number of letters that get sent out saying please come for follow up investigation. When you work that out, it turns out to be just 0.035 or 3.5%. So, it can be incredibly low.

So what's the problem with this? Well, there are a number of problems. I have a book called The Maths of Life and Death - the same times this talk – and in this book I tell the story of Kaz Daniels who is a mother of three from Northampton. Kaz Daniels gets invited for her first ever screen, and a few weeks later she gets sent a letter on the Thursday afternoon saying you need to come for further follow up tests next Monday - so three days' time. She thinks - because of the urgency of the follow-up, because it says you're invited to come for further tests - she assumes, quite naturally, that she has breast cancer. She makes plans about what's going to happen to her kids when she dies, she doesn't eat properly over weekend, she doesn't sleep properly, she suffers severe psychological stress because of this letter. Actually, when she goes to the follow up test on the Monday it turns out that she didn't have breast cancer, which was always likely to be the case, because we know that the vast majority of these positive letters or these requests for follow up investigation or false positives. I think it's important that people understand that there is a high rate of false positives when we have a disease which has a low prevalence in the population, and we are screening lots of people for it.

But that's not the only problem. This is a quote from Muir grey who's the former director of the UK national screening programme. He said “All screening programmes do harm. Some do good as well and of these, some do moregood than harm at reasonable cost.” But he says, “all screening programmes do harm”. What does he mean by this? Well, there is another problem that lots of small, potentially benign, tumours will be picked up with these screens. Tumours which will be so small or slow growing that they would never cause a problem in that woman's lifetime. Nevertheless, when someone hears the word ‘tumour’ or ‘cancer’, something goes off - some alarm goes off in their brain and they think “I have to do everything that I can to avoid this existential threat”. So people often end up undergoing unnecessary surgery, which carries with it its own risks - so with radiotherapy for example, for breast cancer, there's an increased risk of of heart disease for example - so they end up undergoing these dramatic, sometimes life changing surgeries, sometimes full mastectomies, which can be incredibly psychologically difficult to deal with as well as physically - undergoing general anaesthetic, which brings with it its own risks, when actually the tumour that was picked up by this screen maybe didn't have any potential to to cause them damage over the course of their natural lifetime.

The other problem with all these false positives is that these follow up tests are sometimes much more invasive than screen. Screens I gather - I've never been from myself - but I gather they are particularly comfortable, but potentially follow up tests involve things like biopsies which are non-trivial operations to undergo, particularly for things which turn out to be false positives. Screening programmes are not without their problems.

Now the other lesson that I want us to learn about screening programmes is that we should start to expect these false positives. What do I mean by that? Well, if you go for consecutive tests with a false positive rate of 10% - lets assume you get false positives independently on each test – that may or may not be the case with mammograms, there may be a particular reason why you are flagged up as being a false positive or not. But let’s assume that these tests are independent of each other each time you go. If each time there’s a false positive rate of 10% - as there is with mammograms – how many tests do you need to go to before it becomes more likely than not that you're going to have received a false positive test result? We can work this out for a number of tests using a bit of fairly straight forward mathematics.

What I'm going to do is I'm going to say for each number of tests that you undergo, firstly, ‘What's the probability of not receiving the false positive?’ because that's actually easier to work out and then ‘What's the probability of receiving a false positive?”

For one test the probability of not receiving a false positive well that’s just 0.9. OK, so if the probability of receiving a false positive is 10% or 0.1 the probability of not receiving one is 0.9. I'll write that in a slightly circuitous way as 0.9^1, which is just 0.9. So, the probability of a false positive is 10% or 0.1. That's just the false positive that rate that we were given for a single test. It seems a bit trivial for a single test, like I've gone a bit too far in trying to work this out, but it becomes more easy – it makes it easier - when I'm thinking bout more tests than one. When I've gone for my second test, in order to have had no false positives so far on both of these tests I need to have had no false positive on the first test - that happens probability 0.9 - and no false positive on the second test and that happens independently the probability 0.9. So, I have to multiply 0.9 by itself to find out the probability of no false positive being 0.81, and therefore the probability of there being a false positive - or at least one false positive - in those two tests being one minus that at 0.19. So that’s after two tests. After three tests I need to have no false positives on the first and on the second and on the third. So that's 0.9^3 which is 0.729. That's the probability of having no false positive, and so the probability of having at least one false positive during that time is 1 minus that, which is 0.271. So, it's getting higher, you can see, and actually - I'm not going to go through all the calculations - but by the time you get to seven consecutive tests the probability of no false positives in any of those is 0.9^7, which is 0.478, which is just below ½. So, the probability of at least one false positive during that time is 0.522, which is just above ½. After you've been for seven independent screens, or tests with a false positive rate of 10% in each one of them, independent of the others, the probability of receiving at least one false positive outweighs the probability of not receiving a false positive. In short, it becomes more likely for you to have received a false positive than for you not to have received a false positive. It's worth bearing this in mind - that if you go through enough screens then you should start to expect these false positives. For women undergoing screening for breast cancer, for example, if you get invited for screening from the age of 50 every three years until you're 72 or 73 then you might expect to have seven or maybe even eight of these screens over the course of that time period. So, you might expect to receive false positive results.

The take home messages from this part of the talk are:

1. Take screen results with a pinch of salt.
   What I am absolutely not saying is 'don't go for screening'. Please do continue to go for screening. Screening programmes do catch cancers early and it does mean that you can have much better treatment outcomes if you catch cancer early. But what I am saying is that you should, if you get a letter telling you 'you need to come back for further tests', of course you should go for those further tests but, you shouldn't freak out about it. You should bear in mind that false positives are surprisingly likely in these screening tests where we have a low prevalence in the population - so a small number of people have this undiagnosed disease - and the tests that we are using for them lack what's called specificity, so they give a nontrivial rate of false positives. And this isn't just for breast cancer screens, this also goes for things like prostate cancer screens as well - anywhere where the prevalence is low and we're using a test which lacks a little bit of specificity - it has some mental rate of false positives - you need to take your initial screening result with a pinch of salt and go for follow-ups and get it checked out.
2. You should also start to expect these false positives as well.
   The way that I like to think about screening results is: When a company is hiring for a job, they send out an advert and the idea is that they get people to send in their CVs and they can sift through those CVs fairly quickly and easily. That's like a screening programme. We send everyone who was at slightly higher risk of getting this disease, we sent them a letter saying come for a test. We will do a cheap easy-to-use test like a mammogram, for example. It won't be the most accurate test, but we can quickly try and find out the people who are going to need follow up tests. So just like with a job interview you sift through CVs then you say, 'right I'll take these candidates.' This is a broad-brush measure, but what I want to do is throw the more expensive stuff - like the interview which is intensive in terms of the people that takes up or the assessment centre - I'm going to get this person to come down and I'm going to use some more accurate tests to find out whether this person gets the job. That's exactly same screening programme you've got a letter saying, 'come for some more accurate tests so I can diagnose whether you actually have this disease or not' and the whole point of that is - the whole point of the analogy rather is - that just because you get asked for an interview for a job, you shouldn't assume you've got the job. In exactly the same way, just because you've been invited for follow up tests from a screen, you shouldn't assume you have the disease that is being screened for. That's how I would view the results of screening.
3. The other thing is to ask for a second opinion.
   The example at the start shows you that the doctors who are often handling the statistics who are handling our diagnoses and not always the best equipped statistically to understand the mathematics underlying these diseases. I think asking for a second opinion, not just from a different doctor, but also asking for a second test. If you get a second test, even if it's no more accurate in the first test, that can dramatically reduce the rate of false positives. Two tests really can be better than one.

Maths and the Media

I'm going to segue into this second part of the talk, which is about the places where maths can come up in the media - in newspaper headlines, in TV or the radio - but I'm going to continue to talk about medical applications. It seems that almost every day there's a newspaper headline or story that comes up about the impact of one of our lifestyle choices on something to do with our health. So, it might be that we are eating too much of the wrong thing, or we're not eating enough of the right thing, or maybe we're sleeping too long, or we're drinking too much or we're not drinking enough maybe. There's lots of headlines about the about lifestyle choices and the reason they're often in the headlines is because people genuinely want to try to make the best lifestyle choices that they can in order to improve their health and to improve their well-being. It's completely reasonable that people want to do that, but we need to be careful about the way these stories are presented to us. I'm going to tell you a cautionary tale.

Back in 2009 the World Cancer Research Fund published a paper, which was a review paper, which looked at the results of over 100 other papers. It was a 500-page document, and it was reviewing lots of different studies into the impact of various lifestyle choices on the probability of getting a whole range of different types of cancers. The one study that got picked up on by the media was one which was about the probability of getting colorectal cancer if you eat 50 grams of processed meat every day. So, of this 500-page report with over 100 different studies in it, someone at The Sun newspaper had read this or at least read part of it and they have decided to highlight this study about processed meat and that the chances of getting colorectal cancer if you eat 50 grams of processed meat a day. And they decided to sell this story under the admittedly inspired headline 'Careless pork costs lives'.

What they claimed in this this eye-catching headline story was that 'Eating a bacon sandwich everyday increases the risk of colorectal cancer by 20%'. Now I saw this headline and I thought that seems like a huge increase. Could it be that say, the background rate of getting colorectal cancer if you don't eat a bacon sandwich everyday is 5%, could it be that genuinely eating a bacon sandwich everyday increases that risk by 20% to 25%? Do a quarter of all people who eat bacon sandwiches everyday – or 50 grams of processed meat as it wasn't the actual study – to they end up getting colorectal cancer. So, I decided that I was going to investigate the statistics - the truth behind this headline grabbing figure. But, before I had a chance to do that, The Sun had already launched the 'Save Our Bacon' campaign, in which scientists were branded as health Nazis who'd declared a war on bacon. Despite the fact that, of course, the original scientific article said nothing about bacon at all - it said about 50 grammes of processed meat a day. Nevertheless, The Sun realised that pitching this story in terms of bacon – a buttie - would strike at the heart of the British psyche and it would sell papers. And indeed, it did. It sold lots of papers on the back of the 'Save Our Bacon' campaign. So where had this figure of 20% actually come from?

Well, The Sun had calculated something called the relative risk. They'd looked at a number of people in this study. They said, of people that don't eat a bacon sandwich every day, five of those might be expected to get colorectal cancer during their lifetime. And then of the people that do eat a bacon sandwich every day, 6 of these might be expected to get colorectal cancer during their lifetime. And then what they'd done is they'd calculated this thing called the relative risk. Now relative risk is usually calculated for drugs or treatments which are designed to reduce the prevalence or the probability of getting a disease. So usually, the relative risk is the risk with taking the treatment divided by the risk without taking the treatment. In this case the treatment is a relatively pleasant experience, it's eating a bacon sandwich. So, they calculated the relative risk of getting colorectal cancer with eating the bacon sandwich and they divide that by the risk without eating a bacon sandwich. They said six people get colorectal cancer when eat a bacon sandwich everyday and five without eating it, so they figured out that 6 divided by 5 is an increase of – well an increase of one on 5 is an increase of 20%. So, it gives you this value 1.2 and this increase of nought .2 is an increase of 20% on the baseline risk of 1. So, it is indeed a 20% increased relative risk. Of course, in the headlines they didn't mention that this was an increased relative risk but but that's why it caught so many eyes.

What would be more appropriate and more useful for us is not to know just what the increase in the risk is but to know what the absolute risk is. So, in the actual study what they found was, of people in the survey who didn't eat 50 grams of processed meat a day, they found that 5 of those might expect to get colorectal cancer over the course of their lifetime. And when they studied people that did eat at least 50 grams of processed meat a day what they found was that that number increased *dramatically* to… six people in 100, or 6%. So really not a big increase in absolute terms. An increase of 1% for people that eat 50 grams of processed meat everyday and their chances of getting colorectal cancer over their whole lifetime. So, an absolute increase risk of 1%. But of course, a figure 1% doesn't sell many newspapers whereas the relative risk of 20% looks much more dramatic.

So, you should be careful when you just presented with one single headline figure, usually if it's big and if it's a percentage then you need to dig deeper to find out what the absolute risks are there should be two figures and they will usually be smaller and they will give you a better picture of what the chances of getting with disease with and without that lifestyle intervention are.

I'd love to tell you that this was a problem that was just restricted to newspapers and the media - people who perhaps are less scientifically literate, but that's not the case. In fact, you see this sort of problem in scientific literature - even scientists make these sorts of mistakes - and you also see it in patient patient advice literature. This is a screen grab for a tool which was around - you can probably date it from looking at the screen grab - it was around in the early 2000s. It was called the Breast Cancer Risk Assessment Tool and it was developed by the National Cancer Institute in the United States. It was a tool which was designed to give people who'd been diagnosed with breast cancer the information about the risks of various different treatments that were available to them - both the chances of improving their conditions under that treatment but also the potential side effects.

There was one particular drug that was on there, which is called tamoxifen - quite a famous drug, quite well known - which is used for treating breast cancer and this is the way that this tool chose to present the benefits of taking Tamoxifen.

It said, 'Women taking Tamoxifen had 49% fewer diagnosis of breast cancer'. That sounds brilliant. 49% fewer. That's a huge number of fewer cases - it's almost half the number of cases, which is fantastic. But when they presented the side effects of taking Tamoxifen, they did it like this 'The annual rate of uterine cancer was 23 per 10,000 for taking women taking Tamoxifen compared to 9.1 per 10,000 in the placebo arm'. That's really difficult to parse, I think. I think it's really difficult to read, the way that they'd written it. And also, they used, what we saw before are called, the absolute risk. So, they'd used '23 per 10,000' - that sounds like a small figure does 9.1 per 10,000 so maybe this increased risk in getting uterine cancer when taking tamoxifen, maybe that's trivial. Maybe this small number of cases really doesn't make a big difference at least that's what you're meant to believe when reading this. This is called 'mismatched framing' - when you present the relative risk of the benefit or the thing you want to emphasise and the absolute risks for the thing you want to downplay – so the side effects in this place - that's called mismatch framing and it can really make a big difference to the way people perceive that risk.

So, what would be fairer than using this mismatched framing idea would be perhaps to present both of these things using the absolute risks. That would be more informative. For uterine cancer we have already seen that for people taking tamoxifen there were 23 per 10,000 women in this trial that was done on tamoxifen. And for the women they weren't taking tamoxifen in the trial the rate was

9.1 for 10,000. So, an increase for people taking tamoxifen of 14 cases of uterine cancer per 10,000. And for breast cancer, for women in this trial, 34 per 10,000 when taking tamoxifen, which is reduction of about half, as the relative risks suggested from women who were not taking tamoxifen where you get breast cancer at the rate of 68 per 10,000. So again, these are much fairer ways. And actually, you can see that the rates of getting breast cancer and not dramatically higher when given in absolute terms in the risk of getting uterine cancer. They are indeed a bit higher, but not orders of magnitude higher. Alternatively, if you're going to do this - and this is not the best way to present the risk - but if you are going to use the relative risk then you need to present both the side effects and the benefits using the relative risk. So, for breast cancer, we've already seen that tamoxifen gave 49% fewer diagnosis, which sounds like great news. But the reason why the uterine cancer figure - the increase in the rate of uterine cancer figure - wasn't presented using a relative risk is because it looks really really bad. So, this increase - you can calculate it from the absolute risk - this increase of 14 cases from 9.1 to 23 per 10,000 compared to a baseline of 9.1 - 14 is an increase of 153%. So, for uterine cancer the relative risk looks like 153% more diagnosis, which is clearly not something that you want to be presenting, if you are presenting the side effects and trying to sell this drug to people - which is not what the breast cancer risk assessment tools trying to do - but even so this bias, this mismatch framing may have been done sub-consciously. And the irony is that if you'd just presented the absolute figures, you could see that taking Tamoxifen reduces cases of breast cancer by 34 and it only increases cases of uterine cancer by 14, so overall the net saving of 20 cases of cancer for 10,000 shows a significant benefit. You don't need to -at least in this case - you don't need to try and fool people by presenting the risks in a different way. Tamoxifen is a drug which has been used to treat breast cancer for a number of years because the benefits tend to outweigh the risks.

There are lots of different ways of trying to sell statistics. If you've got a number and you want to tell a particular story, then you can use statistics in different ways to highlight different aspects of that number. But sometimes there are some statistics that are so stubborn that you just can't bend them to your will, so what you can do, if you really want to tell a particular story, is what Donald Trump did back in 2015 during the run-off for the Republican candidate nomination for the presidency and that is just to make it up.

This is an info graphic that he retweeted back in November 2015. I want to warn you this part of the talk gets relatively serious. This is an info graphic which purports to show statistics about the ethnicity of people who were killed in United States in 2015 and by whom they were killed: by police or by white people or black people. So, the statistics I want to highlight to you from this infographic that Donald Trump retweeted are these:

It suggested that of all the black people that were killed in 2015, 1% of them were killed by police and of all the white people that were killed, 3% of them were killed by police The other two statistics I want to highlight are that of all the white people that were killed in 2015, apparently only 16% of them were killed by other white people, whereas of all the white people killed in 2015 apparently as high as 81% of them were killed by black people.

Now I should preface this by saying that the source for these statistics is the Crime Statistics Bureau of San Francisco, and as a spoiler alert for what's coming up, I should tell you the Crime Statistics Bureau of San Francisco is completely made up. It doesn't exist and so that's a forerunner of what you're about to see.

I saw this info graphic and I thought I'm going to check these figures out because they don't sound quite right to me. So, I went to the FBI website and I found out that actually the proportion of white people killed by other white people is not 16% it's actually 81% and of all the white people that were killed in 2015 the proportion that were killed by black people is not 81% it's actually 16%. So, these two figures, someone had gone to the trouble of actually looking up the correct figures and then just crudely transposed them to paint a very different picture about who is killing whom in the United States. The other two figures that are highlighted to before in the previous infographic where that, of all the black people killed in 2015 it said that 1% were killed by police. Actually, the figure is as high as 11%. Indeed, of all the white people killed by police, it suggested that that that was as low as 3% in fact it's actually as high as 16%. Now these last two statistics, these come from the Guardian website. The reason for that is because back in 2015 the FBI were not keeping data about the ethnicity of people killed by police officers or law enforcement officers in the United States and the director of the FBI, James Comey, said ‘It's embarrassing that this British newspaper paper is keeping better statistics on this than we are’. And he's absolutely right. So, I had to go to the Guardian and look at their project which is called ‘the counted’, which was actually keeping track of these statistics to dig out those two figures.

So, these are the real statistics. Now this was important because back in 2015 this was sort of the zenith of the first wave of the Black Lives Matter movement. It was incredibly important movement in the United States and still is to this day, but It was reaching its zenith and actually, by about a year later the Black Lives Matter movement had made its way across the Atlantic to the United Kingdom It add started to raise the ire of some right-leaning journalists in the UK, in particular Rod Liddle wrote in The Sun in September 2016. He wrote the following on this very issue, he said: “There’s also no doubt whatsoever that the greatest danger to black people in the USA is . . . er . . . other black people. Black-on-black murders average more than 4000 each year. The number of black men killed by US cops — rightly or wrongly — is little more than 100 each year”. And he finished with this, he said ‘Go on, do the math”. So obviously reading this I couldn't help myself but go and ‘do the math’.

Firstly, a bit of fact checking on Rod Liddle’s statistics. Liddle said that ‘black-on-black murder average more than 4000 each year’. By that he meant that 4000 black people killed by other black people in that year. I looked at 2015, which is the last year that Liddle could conceivably have gathered statistics for, should have wanted to, Actually, in total, the number of killings in which the victim was black were only 2664 so nowhere near this 4000 figure that little was quoting. In fact, of those 2280 were committed by black citizens and 229 by white citizens. So nowhere near the 4000 figure. Liddle also said, ‘the number of black men killed by US cops, rightly or wrongly, is little more than 100 each year’. Well again looking at the Guardian’s ‘Counted’ project, you find out that the real figure is actually three times that amount, it's over 300 black people killed by US cops in 2015, so three times that amount. Firstly, his two statistics were exaggerated: One, to make the number of black citizens killed by other black people much higher and Two, to make the number of black citizens go by law enforcement officers look much lower. For balance, here is the the number of killings in which the victim was white. In total 3167, of which 500 committed by black people 2575 by white people and 584 by law enforcement officers. Of these statistics - and one thing that Liddle really didn't draw anyone's attention to and which I think is important to draw attention to - is that of the total number of black and white people that were killed in the United states, about 45% of them were black and about 55% of them were white, which is which is pretty crazy when you consider that black people account for maybe only about 12% of the population of the United states. Black people being killed with a disproportionately high rate in the United states, but that is not the story that Rod Liddle wanted to tell, but that is part of the story that I want to tell you this evening.

So, what about Liddle's other claim, his last claim? He said, “There is also no doubt whatsoever that the greatest danger to black people in the USA is other black people”. He's using those two figures in the previous slide, the fact that more black people killed by other black people than are killed by law enforcement officers to make this argument that `the greatest danger to black people is other black people not police officers`.

To unpick that and to show you why that's a false argument I want to do a little thought experiment with you. In 2015 toddlers shot and killed 21 U.S. citizens. It's crazy but somehow these toddlers were left alone with guns and they managed to end up killing 21 U.S. citizens. In the same year bears killed just two US citizens. OK so toddlers killing far more people than bears. Then I'm going to ask you the question 'Who would you rather be left alone in a room with? Would it be a toddler or a bear?'

I think the answer is quite clear for everyone. Of course, you'd rather be left alone in a room with a toddler. Toddlers don't kill more people in the United States because they are inherently more murderous or more dangerous than bears. It's because there are shed-loads more toddlers in the United States thank there are bears and people come into contact with toddlers far more often than they do bears. For exactly the same reason, black people don't kill more other black people in the United States because they are inherently more murderous it's just that there are a lot more black people in the United States than there are police officers. So, what you really need to do is to account for the total size of those populations: the population of black citizens in the United States and the population of law enforcement officers in the United States. Then you need to divide the number of killings by the size of the population to calculate what's called the 'per capita rate of killing'. That helps to answer the question "If I'm a black person walking down a dark alley and someone is approaching me and I'm worried that they might kill me, who should be more concerned about it being? Should be more worried that it's another black person or should I be more worried that it's a police officer?' So, let's figure that out. In terms of the mathematics, and we've seen these statistics already. Of all the people in 2015 who were killed who were black, 2280 were killed by other black citizens and 307 by law enforcement officers. Now, when you take into account the sizes of these two populations, there are over 40 million black people in the United States and I just 635,000 law enforcement officers. To calculate the 'per capita killing rate' you need to divide this number of killings by the size of the population doing the killing. It turns out that for black people the rate is just about one in 17,000, whereas for law enforcement officers it's way, way higher. It's one in 2000 almost, so much, much higher than the rate of killing by black people on other black people. Way way higher flow enforcement officers, which really tells a very different story to the story that Rod Liddle was trying to tell us.

Now there are some caveats to this in that, of course law enforcement officers are routinely armed they're also, often in the course of their work, going into conflict situations - situations where perhaps, in the United States, gunshots are more likely to be fired. So, there are caveats for why this rate might be higher. Nevertheless the per capita way of killing for each member of this population is way higher for the enforcement officers than it is still black people. So, it gives the lie to this idea that black people should be more scared of other black people than they are of law enforcement officers in the United States, and it really changes the point of view that Rod Liddle was trying to put across.

The take home messages from this part of the talk:

Firstly, when it comes to medicine, when it comes to reading newspapers, when it comes to listening to the radio and TV, we need to look out for these relative risks. If you're just given a single figure, usually a percentage and often if it's large, then you need to be aware that that might be a relative risk. If you really want to know comparison between two different treatments or two different lifestyle choices then what you really need to know is what the risk is with doing that treatment and without or what the risk with doing that lifestyle choice and without doing that lifestyle choice, if you want to make a fair comparison and understand the magnitude of the size of that risk.

The other thing, of course, is that you should be aware of mismatched framing. If you see one statistic painted as a relative risk and the others downplayed as an absolute risk, then you should of course be wary of what this story is trying to tell you - what this newspaper story or what you're trying to be sold by this company who is telling the story in this particular way. So be aware of this mismatched way that things can be framed.

The final piece of advice from from this, which is probably good advice for everyone in every country, is don't trust politicians because they lie out of their teeth all the time as part of their job! OK, so less of a mathematical message more of a philosophical message there!

Maths of Crime

In this last segment, I'd actually already segued into the mathematics of crime looking at those those killing statistics from the United States. I'm going to continue about theme, but I'm going to take a little mathematical diversion too to build up some maths that I'm going to need to look at particular problem that I'm going to present later in this section. I'm going to start by introducing a relatively famous mathematical problem called 'the birthday problem' If I was doing this talk in the room with you then I'd be saying 'What's the probability that two people in this room share a birthday?" Unfortunately, we're doing this virtually, so I'll rephrase that question and I'll say, perhaps put it another way, "How many people do I need to have in a room before the probability of two people having the same birthday becomes more than ½?" – Before it becomes more likely than not that two people share a birthday. Now I think when most people think about this problem, without thinking too hard about it mathematically, the number that comes to mind is probably 180 or so, because that's about half of 365 - the number of days in the year. So perhaps 180 people will be enough to make sure that the probability of two people having the same birthday becomes more than half - becomes more likely than not?

Actually, the real answer - and the reason why this is a famous mathematical problem is the real answer is really surprising - in fact you only need 23 people in the room for the probability that two of them sharing a birthday becomes more likely than not. Now that seems like a surprisingly low number, so where does that figure, where does that 23-number come from. Well, I'll try and introduced it to you in in the in the following way. Really what's important when we think about this birthday problem is not how many people there are in the room, but actually how many pairs of people there are in the room, because we're talking about two people sharing their birthday that's something that pairs of people do together. So, I want to work out a formula for how many pairs of people there are in the room when I've got a certain number of people in there. And I'm going to start by doing just five people, because I know that I can check that relatively easily. So, I've got five people here - my maths reservoir dogs – I've got Mr. Green, Mr Orange, Mr Purple, Mr Blue and Mr White. I'm going to figure out how many pairs there are. To do that I'm going to get them to shake hands with each other. You can tell that I developed this idea before the times of COVID because now should be fist bumping or they should be elbow bumping or they should probably be standing across the room and nodding at each other or maybe even just outside somewhere that's well ventilated and socially distanced. Anyway, let's pretend that they can shake hands with each other. So, here's Mr Green and he shakes hands with Mr orange and Mr purple and Mr blue and Mr White and that's four handshakes for the four other people in the room. Then Mr orange steps up and he can shake hands with Mr Purple and Mr blue and Mr White, and that's three more handshakes added to the total. Then Mr purple steps up to the mark and he can shake hands with Mr blue Mr White and that's two more handshakes. The last handshake that hasn't happened yet is between Mr blue and Mr. White and that's one more handshake. So, it's no coincidence that with five people in the room I start at 4 and add up three and two in one. I add up the consecutive whole numbers going up to one below the number of people in the room - so four in this case. When I add those up it's quite simple it's just 10 people in the room. So, I know there are 10 handshakes or 10 pairs of people when I've got five people in the room, but I want to generalise this to make it more general - so what happens when I've got more than five people in the room? Well, I'm going to take another diversion in order to be able to answer this. These numbers when you add up the consecutive whole numbers, they're called triangular numbers. Why are they called triangular numbers? Well let's see. Hopefully, if you were brought up in Britain, then you will know the song the 12 days of Christmas. If you weren't brought up in Britain, then it's a crazy song. It's about boyfriend or girlfriend you've got – a partner, your true love - they send you presents after Christmas - for each of the 12 days after

Christmas they send you presents. On the first day they send you a partridge in a pear tree. Let's assume that partridge in a pear tree is just one present for the sake of argument. On the second day they send you the partridge in the pear tree and the two turtle doves. That's three presents on the second day. On the third day they send you the partridge in the pear tree and two more turtle doves and three French hens. Then on the 4th day they send you the partridge in the pear tree, the two turtle doves, the three French hands and four calling birds. On the fifth they send you all those presents once again plus 5 gold rings – that's how you sing it in the song, but that's when you get to 5. That's the reason why they called triangle numbers, because you can arrange them in this nice triangular shaped array. People who are interested in snooker will be familiar with this because when you make the triangle of reds you make a nice triangular shape with it and that's one red ball at the top, two in the next row, three of the next, four and five. OK so it's the same number as this this. So, the question we're going to ask is "How many presents did my true love sent to me on the 5th day of Christmas according to the song?" How am I going to calculate this? Well let's start by rearranging these objects. Let's start by making them into a different sort of triangle - a right angled triangle. Then I'm gonna bring a whole second set of these objects to make another different shape - in this case a rectangle. Now I've got a rectangle of objects, which is going to be easy to calculate. I just need to calculate or multiply the number of rows - 5 - by number of columns – six. OK so I've got 5 rows, because it was the 5th day of Christmas and I've got 6 columns because there's one extra when I've added this on. Generally, I'm going to be calculating n - where n is for the n^{th} day of Christmas - I'm going to calculating n for the rows times n+1 for the columns. That's going to be the number of objects in this rectangle, but I've got to remember to divide by two, because I've added in twice as many of these objects as I should actually have. OK, so the formula for the number of objects that I would get on the end day of Christmas is $n(n+1)/2$. OK, so when I plug that in for five that's 5 * 6 / 2 which is 15 presents, which people who who play or watch snooker will be familiar with - the fact that there are 15 red balls in this array.

That's good that's my formula for the number of presents on a particular day of Christmas and it generalises. How does that work out a handshakes? Well, with 23 pairs of people in the room, the first person can get up and shake hands with 22 other people, the next person with 21, the next person with 20, the next person with 19 – and I don't know why I animated the whole of this triangle because, of course, I've now got a formula to work this out anyway. I need to work out the sum of all the consecutive integers from 1 all the way up to 22. I've got this formula that – it's going to be 22 times 23 – so $n(n+1)/2$ and that's going to be 253. So, when I've got 23 people in the room, I've actually got 253 pairs of people, which goes some way to explaining why it's so much more likely to have a pair of people that share the same birthday in the room than you would expect with just 23 people in the room.

But I'm not done yet. So, I figured out that with 23 people in the room there are 253 pairs of people, but actually want to do the calculation that tells me what the probability of two people sharing a birthday is. How do I get to that? Well, firstly I'm going to work out what the probability, p, of 1 pair of people not sharing birthday is. The way to work this out is a bit like what we did with the mammograms earlier. When working out the probability of false positives it was easier to work out the probability of not having a false positive first and then working out what the probability of getting a false positive - or at least one false positive - was by subtracting that value from one. So, I'm going to do exactly the same thing here. I'm going to start with just one pair of people. What's the probability that one pair of people don't share a birthday? Well, the first person can have their birthday on any day of the year. I put this on here as that's my birthday April the 4^{th}. The other person in the pair can have their birthday on this day or this day or this day – don't worry I haven't animated them all – or on this day, or in this day, or on this day and so on. In fact, that person can have their birthday on any of the remaining 364 days of the year. So, the probability that when they were just two people in the room - just 1 pair of people - the probability that they don't share a birthday is incredibly likely. The probability they do share birthdays really unlikely - it's 1 divided by 365. That's where we're starting off. Now what happens if I take that probability and I try to figure out what the probability, P, the 253 pairs of people don't share a birthday. OK, so I have to have to think, ‘Well if I need one pair of people not to share a birthday, but I also need the next pair not to share a birthday and the next pair on the next pair… If those birthdays all independent then I need all 253 pairs of people not to share a birthday, which means I have to multiply that probability of one pair of people not sharing a birthday - 364 over 365 - by itself 252 more times, or equivalently I raised it to the power of 253. Now this number 364 divided by 365 is pretty close to 1. But actually, even when you multiply a number - even in number that is that close to 1 - by itself enough times that number starts to get smaller and smaller - the product starts to get smaller and smaller. Actually, by the time I've done it 252 times, the probability that two people in a room of 23 don't share a birthday is 0.4995 so just below ½. So that means the probability that two people - at least two people - do share a birthday, in a room with 23 people in, is just above ½ - its one minus that figure.

OK, so for different numbers of people in the room I've plotted the way that the probability changes with that number of people. I have got the number of people on the horizontal the axis at the bottom there. On the side - the Y axis, on the vertical axis - I've got the probability of seeing a match. So, what you can see is, as the number of people in the room increases, the probability of having a match increase. I stopped it here at 23 people to show that the probability of having a match becomes bigger than 1/2 when you have 23 people in the room. You can see, by the time that I've got up to maybe 60 or 70 people in the room that the probability of two people sharing birthday is almost certain. When I’m doing this talk with people in person, I often get people to try and make a bet with me. I will often put a pound on and say “I bet two people in this room share a birthday” - even when there are small numbers of people in the room as long as there are more than 23 - and see if someone will take the corresponding other side of the bet. People always do because they think it's so unlikely that two people will share a birthday. In fact, even if there are people not willing to take that one-to-one odds, as long as I've got enough people - maybe about 60 people in the room - I could offer odds of 10-to-1 and still expect to make money on that bet. It's a surprising result and they think it's something that people really find difficult to intuit.

I've done some theory with you. We’ve done a little bit of mathematics, but what about the actual real world application of this? Does it actually work in reality? Well back in 2014, the World Cup was hosted in Brazil and the great thing about the World Cup in 2014 was, apart from various of the games, the really great thing for a mathematician was that each of the 32 squads had 23 players in it - this magical number where we would expect the probability of two people to share a birthday in that group of 23 to be more likely than not, but only just - it's about a half. This is the Brazil squad of 23 people and lo and behold you've got two players - Hulk and Paulinho - who share a birthday on the 25th of July. So, it works in this case, but that's just one example that i could have cherry picked. There were 32 teams in the world cup in Brazil in 2014. The great thing is that, if you're a mathematician and you are so inclined, then you can look up the birthdays of all the 32 * 23 people who played in that world cup or who were part of squads in that World Cup and you can figure out ‘of all those 32 teams, what proportion of them had two players - at least two players - that shared a birthday?’. If you chose to do that, as a mathematician might have done, then you find that exactly half of the squads - so 16 out of 32 – have two players, or at least two players that share a birthday. Now that's what this probability would have predicted - that would be your best estimate of the number of teams. Of course, it doesn't always work out to be perfect – that it is going to be exactly a half of these teams. It is probability, it’s random, so you get more or less or some occasions but the fact it works out at 16 out of 32 is so beautiful and it really demonstrates that this formula does indeed work in reality.

I'm just moving on to the last part of the talk now and I want to warn you this part of the talk is a little bit more serious and potentially upsetting. But I think it's important to talk about it.

Back in May 2017, a few days after this is the scene in St. Anne’s square in my hometown of Manchester. The reason all these floral tributes and balloons were laid out in St. Anne’s Square was because on the 22nd of May 2017, at about 10 o'clock, 10:30 in the evening and Ariana Grande concert was spilling out. Young people who'd been at the concept were in the foyer meeting up with their parents and people who'd come to pick them up and heading home. At the same time Salman Abedi walked in alone wearing a rucksack full of explosives and bolts and nails and screws and he blew himself up and he killed 22 innocent people that evening in Manchester. People who were young and in the prime of their life and he took that away from them. As I mentioned, my hometown was Manchester and I'd been to concerts in this arena so I was particularly interested in following this because it struck home that this was in my hometown this felt like it could have been me. At the time it wasn't in Manchester, I wasn't even in the UK. I was in Mexico and because of the time delay in Mexico we were in the middle of the afternoon when this was happening. I followed this avidly throughout the rest of that day as the news stories came in and actually, probably most of the rest of the people in the United Kingdom were still asleep, only to wake up to this horrible story the next day. I followed it with great attention over the next few days precisely because it was in my hometown of Manchester and and it struck a really struck home to me. What I notice over the next few days was something interesting about the date - the 22nd of May. I noticed that exactly 4 years before, in another Islamist terrorist attack, fusilier Lee Rigby was brutally murdered outside his barracks in Woolwich. I thought 'That is a strange coincidence for these two terrorist attacks to have occurred within four years of each other.' I wasn't the only person to have notice this coincidence. The Daily Star ran a head-line which said 'Dates mattered to Jihadi terrorists. Manchester arena attack on Lee Rigby anniversary. In that story, it was basically based on a tweet by Sebastian Gorka who was then President Trump's deputy assistant. He said, 'Manchester explosion happens on 4th anniversary of the public murder a Fusilier Lee Rigby.' He said, he concluded from that 'Dates matter to Jihadi terrorists.'

I read that tweet and I thought, 'Is it true? Is it true that Jihadi terrorists are really organised and can strike at will, meaning that potentially we should be more scared of them? Or is it really the case that actually, they don't really know what they're doing and they're just opportunists that take any opportunity that comes along, and they don't really have this great organisational power so that they can strike at will?' So, I wanted to figure out, 'If terrorist incidents were just happening purely at random what's the probability that two of them would occur on the same day?' Fortunately, we have just that information from the birthday problem. 'Are terrorists really that coordinated that they can strike at will?'

Well, I went up and looked up all the terrorist incidents that had occurred against Western nations - Islamist terrorist instances I should say - against Western nations between April 2013 and April 2018. It turned out that there were 39 of these terrorist attacks. Then went back to that graph that I showed you earlier and I plotted this on. I said 'Well, with 39 terrorist incidents, even if they were all occurring completely at random on independent days of the year - random days of the year - the probability that two of them would occur on the same day of the year purely by chance is nearly 90%.' in short we will be, or we should be, far more surprised if two of them had occurred on the same day of the year then when two of them eventually did end up occurring on the same day of the year. It puts the lie to this idea from Sebastian Gorka that somehow just because these two terrorist incidents occurred on the same day of the year that that somehow means that these terrorists are organised. He was tweeting, drawing this conclusion, trying to further his anti-Islam agenda in the United States, making us more scared of these terrorists than we should really be. Doing the terrorist job for them, in short.

I want to finish with these take home messages.

Firstly, this last story illustrates that coincidences, although they seem unlikely, they are surprisingly likely to occur. Given enough time and enough chances and enough events, coincidences can be surprisingly likely but, as humans, we're great at spotting these patterns and then our next thing to do is to draw a causal inference based on spotting a coincidence. Actually, that's not always the best thing for us to do because coincidences do just happen. You might walk down the street and see someone you haven't seen for 10 years but it doesn't mean that somehow magically there's some meaning behind that (usually at least there isn't). Coincidence this can be surprisingly likely. Another message that I want to come across, not just from this part of the talk for the whole talk, is that we shouldn't be blinded by the illusion of certainty. The idea that people that wield the numbers are always in possession of the whole truth. It might be you who is convicted incorrectly by a doctor, who is it was a medical expert but not an expert in statistics, yet they came up with statistics which which gets you convicted. This is a sort of story which comes up all the time in The Maths of Life and Death - in my book. You might be given the wrong information by your doctor because, despite the fact they are fantastic practitioner maybe they're not on top of the statistics. You might read something in the newspaper and change your lifestyle because of that thing you read – because of the increased risk you read about, but actually that risk was a relative risk not an absolute risk and actually the real danger was massively over-amplified.

When you see statistics out there in the wild, remember that statistics can be manipulated to tell a particular story. Numbers are not always these nuggets of hard objective unquestionable truth that we like to think they are. We need to question the numbers and question the people who might wield those numbers against us.

The last message from the talk, as I mentioned I often use this talk to make a bet with people - to try and fleece money out of innocent attendees of this talk - so I would say if a mathematician offers you a bet, don't take it. So, beware of mathematicians bearing gifts.

That's where I'll finish, and I'll say thank you to you for listening. This has been absolute pleasure to give this Gresham lecture and I hope to see you again soon.

Further Reading

The Maths of Life and Death, Kit Yates, Quercus 2019 - https://www.amazon.co.uk/Maths-Life-Death-Kit-Yates/dp/1787475425