



© 1997–2009, *Millennium Mathematics Project, University of Cambridge.*

Permission is granted to print and copy this page on paper for non-commercial use. For other uses, including electronic redistribution, please contact us.

January 2005

Features



All about averages

by Andrew Stickland



Mean or median?



Another rainy day in London

All about averages

I'm just off to meet my cousin Jean-Christophe who's visiting from Paris. It's raining, but as J-C has pointed out to me on numerous occasions, it's nearly always raining in London. He's always telling me how lovely Paris is and how it hardly ever rains there. As he hates the rain, he wonders why he ever bothers to come over to London at all.

And he's right. Paris has more dry days each year than London – a lot more. But I've just been checking my international weather database and I've got some interesting figures to show him. You see, the average annual rainfall in Paris is nearly three times as high as in London.



Where it's sunny every day

So by coming to London, J-C is actually getting away from the rain, even though it looks like it's going to rain every day while he's here.

It sounds like a paradox – Paris has almost three times as much rain as London but London is much rainier than Paris. The answer is simple. In Paris when it rains, it pours – unlike the constant drizzle we're used to over here.

This is a nice example of how averages can be misleading. The fact is, for someone wanting a dry holiday, what matters is not how many inches of rain will fall when it rains but how long it's likely to rain for.

Anyway, the rain has eased off a bit and I've persuaded J-C to come out, and here we are sitting drinking coffee and eating croissants like true Parisians. And, interestingly, he wants to talk about averages.

Yesterday he read an article in the local paper which said that the average commute for people living on my street was 50 miles a day and he's wondering how we can all bear such long journeys to work.

I'm shocked by this. I know most of the people who live on my street, and they're nurses, teachers and office workers – they all work just around the corner. How can the average possibly be that high? Then I remembered. There's a woman a few doors down who goes out to New Zealand four times a year to her head office – notching up about 95,000 miles a year: that would certainly up the average.

All about averages



Commuter vehicle

But this is stupid. One person travels round the world four times a year, and practically everyone else walks to work. So what use is knowing the average distance travelled?

So in this case we need a better way to work out the average. Instead of adding together all the distances travelled and then dividing by the number of people in the street, we'll imagine lining everyone up, starting with the person who has the shortest journey, and ending with the woman who travels to New Zealand.



How many rainy days?

The person in the middle is then pretty typical – half the people on the street travel further and half travel less far. Now we have a fairer indication of how far a typical person in my area travels to work – and it turns out to be about a mile. (This sort of average is called the *median* – the other sort, where you add up all the distances and divide by the total number of commuters, is called the mean.)

So which method is right: the one which give us 1 mile, or the one which gives us 50?

The trick is to know beforehand exactly what information you need. Umbrella sellers don't care what the annual rainfall is, they care how many rainy days there will be. A specialist in flood defences, however, needs to know how heavy the rain will be when it comes.

All about averages

Here are two more examples of how knowing an average is not always as helpful as you might expect. The average size of a household in the UK is currently 2.4 people. But how many households of 2.4 people do you know? Or how about this: over 99% of the population have more than the average number of legs.

Simpson's Paradox

It's Saturday afternoon and I'm exhausted. I've just been forced to spend the whole day shopping for clothes. I've spent a small fortune and I've collected enough bags full of clothes to keep me going for ... well, until I'm retired, probably!

It's all because of a bet with J-C. He says I'm not very careful with my money. He says I spend much more than I need to on clothes and I should be more careful with my purchases ... like him. But I'm sure he can't possibly spend less than me on clothes. After all, he's French, he dresses well. He must be extravagant.



How much per item?

So we've made a bet. We've gone out on a shopping trip together, both buying what we want, and we'll see who's the most extravagant. We agreed that the best way to decide was for each of us to add up how much we spent and divide the total by the number of items we bought. That way we should get an average cost per item, which seems quite a fair way to decide.

In the first shop we went into, I bought 5 items and spent £200, so my average was £40 per item. J-C only bought 3 items, but only spent £90 to get them, so his average was £30 per item. £10 less per item than me. So I'd better start shopping more carefully. In the second shop, I only bought 2 items, but I had to spend £160 to get them. So my average was £80 per item. J-C bought 6 items costing a total of £420. So his average was £70 per item.

Again, he ended up spending exactly £10 less than me per item. So, according to J-C, I'm not a careful shopper, I've lost the bet and I have to buy dinner.

But wait a minute. What about the total for the whole day? I bought 7 items and spent £360; J-C bought 9 items for £525. My average is $£360/7=£51.40$. J-C's average is $£525/9=£58.33$. It seems J-C has spent over £5 more per item than me. Hmmm, dinner's on him, I believe!

All about averages



And how much each for these?

So what just happened? In both shops, J-C spent less per item than I did, but when we take the average over the whole day, I've ended up spending less per item. How can that be? It seems like a paradox.

And it is. It's called *Simpson's Paradox* and it's a well-known puzzle for budding statisticians to think about. To put it technically, Simpson's Paradox occurs when the direction of an association between two variables is reversed when a third variable is controlled.

In our example, the two variables are the number of items bought and the total amount of money spent. The third, or controlled, variable is which shop we were in at the time.

But why does it happen? Well, it may help if I tell you the names of the shops we visited. First stop – where I bought 5 things and J-C bought 3 – was Marks and Spencer, well known for its reasonable prices. Then we decided to splash out a little and we ended up in Emporio Armani, where the prices were a bit steep for me and I only bought 2 things, whereas J-C, possibly confused by the exchange rate, bought 6.

So, although in each shop J-C bought cheaper things than me, he did more of his shopping in the more expensive shop. Mystery solved.

The message here is that you can't average averages. Sometimes they just don't work the way you expect them to. But Simpson's Paradox is not just a party trick. It does have an important part to play in our understanding of statistics and averages – or at least it should do.



He'll be earning more – but how much more?

All about averages

Here's a serious example. On average, a working woman only earns about 80% per hour what a man does. Obviously this isn't right. Two people doing the same job should be paid about the same. But in fact they are. Most of the gap is accounted for by the fact that more women work in the lowest paid jobs – the equivalent of my doing more shopping in Marks and Spencer – but most of the really high earners are men – the equivalent of J-C doing more shopping in Emporio Armani. The gaps between male and female earnings are much less if you make your comparison job by job.

And there are many other instances. For example, drug trials which show that a new drug is better than an existing one for both men and women separately, but less effective than the existing one when the two groups are combined. Or the results which seem to show that too few female applicants are being accepted onto university courses, despite the fact that for each separate course, the acceptance rate is the same for both sexes, or sometimes even higher for women than for men (more women than men apply for the most popular courses; more men apply for the less popular ones).



Which is better?

These are real issues, not just thought experiments, so, remember – averaging averages is not always a good idea.

Outliers



The way the neighbourhood is going?

I have some interesting news for you. I'm turning into a criminal – a burglar, to be precise. You may well find this fact rather shocking, but then so did I when I first discovered it. I'm not planning to break into any houses

All about averages

in the near future, and I've never broken into any in the past – except for the time I had to break into my own house at midnight on a Saturday night, but that was because I'd lost my keys.

But I have the figures right here in front of me. Last year there was only one burglary committed in my neighbourhood, but this year there have been 300. If we assume that my neighbourhood is made up of 1000 people, that means that on average, 3 out of every 10 people that I pass on the street are the victims of burglaries.

Which is about right if we also assume that each person has only been burgled once. But more alarming is the fact that, on average, 3 out of the same 10 people are burglars. Suddenly I live in an area where more than half the population are either criminals, or the victims of crime. That's terrible. I think I should move house.

But there's some flawed logic here, which you may already have spotted. While it's not so unreasonable to assume that each person has only been burgled once, it is certainly not reasonable to assume that each burglar commits one, and only one, break-in. Many burglars commit dozens, or even hundreds, of crimes before they're caught.



Meet the new neighbour

Here's a possible scenario. Someone new has moved into the area – let's call him Bill – and he's responsible for the entire increase in burglary. This isn't at all implausible – police figures indicate that in some areas practically all theft and robbery are committed by a handful of people, usually to pay for drug addiction.

So the "average" criminality we just calculated is in fact totally unrepresentative. How about trying the other way we saw of finding an average? If we line everyone up in order of number of crimes committed, and take the person in the middle, we get the median number of crimes committed by each person.

But there's a problem here too. There are 1000 people lined up, but only the very last person – Bill – has committed any break-ins. So doing this would give an average of 0. This is great news for the area – but hardly representative of the experience of victims of crime.

The problem is that Bill is what we call an *outlier*, that is, a single result that is so extreme it has an unduly large effect on the average. If we calculate the average one way he makes almost a third of the neighbourhood look like criminals, but if we calculate it the other he disappears.

So what can we do about this? Well, the short answer is nothing. It's simply one of the problems with averages. Sometimes they just don't work the way you want them to. Which is all very well, but it's not much help if you really need to know what a particular average is.

Outliers

All about averages

A better answer is perhaps that you shouldn't look for quick results. If you suspect that your results could be adversely affected by outliers, try thinking of a different way of obtaining them. Perhaps you need a much larger data set first, so that an outlier might not be so extreme. Or perhaps reducing things to a single number is not the best idea.

Averages are tricky things. They're not just about manipulating numbers according to some recipe – you have to understand the numbers you have, how they were obtained, and even why they were obtained. Are they the numbers you need? And what are you going to do with the average once you have it?



Bill's new neighbourhood

One final thought: If Bill is ever caught and sent to prison, he may well find that his new neighbours have far more experience at burglary than he does. If that's the case, by going to prison he will have succeeded in reducing the average criminality of both his old home and his new one!

About the author

Andrew Stickland with his son, James

Andrew Stickland is a freelance writer and copyeditor. He has written articles for various publications on subjects as diverse as travel and war-gaming, and regularly reviews some of the most accessible mathematical books and films, as well as providing the occasional photo for *Plus*.

About the author

All about averages

The genesis of this article was in three pieces, co-written by Andrew and Helen Joyce, the editor of *Plus*, for the Radio 4 programme, *More or less*. They are currently airing, and can be listened to on the [More or less](#) website. Andrew and Helen would like to thank Michael Blastland, the producer of *More or less*, for his helpful comments and input.



Plus is part of the family of activities in the Millennium Mathematics Project, which also includes the [NRICH](#) and [MOTIVATE](#) sites.