Founder's Lecture Transcript

Thank you Maggie, and thank you all for being here. It has been a great treat to work alongside Maggie for these last, nearly five years. You may have noticed that there is theme running through some of the things that I have been involved with, which is aging, something which feels more and more close to home as the years go by, and I realise as I was thinking yesterday afternoon about this that it is in fact very nearly forty years since I first set foot in the College, when I came to be interviewed for a place here in, well I know the date, it was October 11th 1977. I came up on the train from Swansea where I was at school, and, to give you a sense of what it was like in those days, the naivety that was feasible, there was one boy from the school that I was at who'd ever come to Oxford or Cambridge, and he was in his second year, reading Law at Balliol at that point, and he'd arranged to take me out for lunch. So he took me out to Brown's, which was at that time very new, and we started eating. About half way through I turned to Tim, and I said 'Tim, I think somebody's been sick', and Tim said 'No Andrew, that's parmesan cheese'. First time I'd ever come across the smell of parmesan cheese! That was one way in which coming to Oxford changed my life, but more seriously, it has transformed it, and I feel an extraordinary debt of gratitude, I expect many people here do, to the university, but especially to this particular institution. I could not have been more lucky than with the friends that I made here, and also the people who taught me. Maggie has already mentioned John Kay, many of you will know John, an exceptional man, a man of enormous intellectual ability, but also extraordinary kindness, and he has throughout my life been, both a model for me, and someone to try to imitate, but also who has been generous spirited, and was generous spirited then in way that captures the whole essence of what institutions like this are trying to do. They're trying to make the opportunity for intellectual excitement available to everybody, and do it in a way that is I think astonishingly effective. So I think John, but others who working here and teaching then, people like David Bevan, Leslie Macfarlane, Gordon Baker (the late Leslie Macfarlane, the late Gordon Baker), Ross McKibbin, Andrew Oswald who was here briefly, Peter Hacker, you know, wonderful people who gave an enormous amounts of time to youngsters who didn't always get quite as far as they might have done intellectually.

But one of the things that that whole thing left me with is something about which I'm absolutely unashamed, and that is a love of numbers. Numbers, statistics, they really, really, really, matter, and our culture is still one where that's not accepted. That it can be sort of ok for a very senior politician, responsible for looking after the economy, to say jokingly, that he's not very good with numbers. Well that's either a lie, or a scandal. We won't identify the particular individual about whom I'm talking, but I suspect the reason that he did that was that, it was a lie, this particular person was very good with numbers, but felt that it was sort of 'cool' not to be good with numbers. Well that supposition, which I could not

deprecate more severely, it's absolutely scandalous. Very few of us, I suspect, even amongst the most mathematically pure, would stand up and jokingly say, 'I don't read very well'. We think that reading is quite important, that it's a central part of being a civilised and analytically capable person. Well, some kind of facility with numbers, not necessarily mathematics, but facility with numbers, is essential. Without numbers, without a decent understanding of what the world looks like, we can't understand what's going on in the world. Numbers and statistics at their best force us to be precise in thinking, which is one of the things I hope I learned to try to do, even if not to do, when I was here, and then finally, if they're misused, which they often are, then they can do enormous damage.

Now I'm going to begin by asking you a multiple choice question. Some of you may have been subjected to this kind of thing before. I keep the data, although it's anonymous. You're going to have to answer by raising your hands, and, I've been playing this game for thirty years or so, and there's only group that has ever refused to answer by raising their hands and insisted on answering anonymously on little bits of paper; and it was a very good thing that they did that, because their performance was particularly poor, and that was all of the permanent secretaries at Sunningdale, in, probably about 15 years ago, and I'll tell you later some of the things that they didn't quite get right. So this is a very simple question. Imagine that we take the whole population of the UK, the household population of the UK, and we line them up in order of their income. So we put the highest income people on the side of the President of St John's, and the lowest income people on the other side to the President of St John's, and they're lined up in order of income. This is net income, after tax has been paid and social security benefits have been received. Then we take the person 10% from the top, the person at the 90^{th} percentile. These are household incomes adjusted for household size, and we'll assume every household is a childless couple, which might for some of you be an ambition, either the childless bit or the couple. So we take the person who's 90% of the way up the income distribution, and then we look at the person 10% from the bottom, and we divide those two incomes one by the other. So we divide the income of the person 90% from the bottom by the person 10% from the bottom. We ask the simple question, what's the answer? Is it 4, 6, 8, 10, or 12. So it's a pretty straightforward question, and you'd think it was quite an important question to know the answer to if you're going to talk about what's going on in this society. So, so that nobody can't accuse me of bias we'll start in the middle. So raise your hand now if you think the answer is eight. I should also say that I've done this often enough that I'm very good at identifying people who don't vote at all, and if you don't vote at all, then you have to come and stand at the front and sing a hymn of my choosing. So unless, actually, I can see one or two people who might think that's fun, so, if Mr Whittaker doesn't answer, he'll have a different forfeit. Basically, you'll have to do something you don't like if you don't answer. So what's the ratio of the income at the 90^{th} percentile to the income at the 10^{th} percentile, is it 4, 6, 8, 10, or 12? Who thinks the answer is 8? So this kind of

shows that psychology is probably more interesting than economics. I think just two people, with groups of fewer than about 70 [member of the crowd interjects 'can I just ask again what it is you're asking us? Are you asking us how many times bigger it is?] Yes, what's the ratio? So, what's the 90^{th} divided by the 10^{th}: 4, 6, 8, 10, or 12. So two people thought the answer was eight. Who thinks that the answer is 6? Who thinks that the answer is 10? Who thinks that the answer is 4? Who thinks the answer is 12? Well, my estimation is that 5% thought that the answer was 4, 20% thought the answer was 6, 2% thought the answer was 8, 50% thought the answer was 10, and about 15% thought that the answer was 12, which means I've got 7% missing, but I'm willing to blame my estimation rather than people failing to vote. I'll tell you the answer later. Now, the President and I have had a sleep, but I confessed to Maggie earlier that I have on occasion fallen asleep in seminars, and so I know the temptation to which you're all subjected at the moment, and I want to keep you on the edge of your seats. I'm going to show you my favourite chart. This is my favourite chart. I really, really love this chart, I love my wife and children more, but I love this chart a lot. So this is a picture of the income distribution in the UK in 2013. So up here we've got number of people, and along here we've got pounds per week. So to get to annual we've got the multiple this number by 52, which will keep you awake if you're not paying attention. This is absolutely crucial. So to be serious just for a moment, if you're thinking about social policy, taxation policy, any form of government intervention, you need to know what shape this diagram is. And it's quite surprising. So the middle of the income distribution, the median income, median household net income in this country, in 2013, was £458. That's a bit less than £24,000 a year. So half of the households in the UK live with net incomes, adjusted for family size for a childless two adult couple, of £24,000 a year. The bottom 10% here, is at £241 a week, so that's about £12,000 a year. So the poorest or lowest income 10% of the UK income distribution have incomes of less than twelve and a half thousand pounds. And the top 90%, the top 10%, is at £921 a week, or a little bit less than £48,000 a year. So, in 2013, if your household net incomes adjusted for size was more than £48,000 you were in the top 10% of the UK income distribution. So my guess is we have one or two people in this room in the top 10% of the income distribution. If you were to look at the political discourse about 'middle Britain', you might be forgiven for thinking that middle Britain is about here [gestures higher up the chart]. Quite often you'll have middle Britain described as a GP married to a teacher. Well a GP married to a teacher is about here [gestures beyond the board]. Now, at one level it's quite funny that people don't know what the shape of the UK income distribution is, but on another it's not funny at all, it's deadly, deadly serious, because if you're going to be interested in how we should pay for the health service, how we should run an education system, what the tax system should be like, what we should do for the incomes of older people, then you need to know the shape of this picture. And one of the things that John Kay, my tutor and mentor instilled in me very early on working at the IFS in the

early 1980s, was the importance of knowing this kind of stuff. So at the IFS, under John's leadership we spent an enormous amount of time trying to pin down the answers to these kinds of questions. So this is the chart for 2013, which is interesting on its own, but what would be really, really interesting obviously would be a consistent time series of this cross sectional distribution going all the way back for as far as we could go. This data comes from the Family Resources Survey. It used to come from a survey called the Family Expenditure Survey, which was first produced in 1961. It was mainly produced to get the weights for the retail price index, but it also contains income data. When I started working at the IFS in 1980 we were working on the 1977 Family expenditure survey, which had 7198 households in it, three of which were corrupt (their data was corrupt not their behaviour), and it took an enormous amount of time and energy to create a chart like this for that one year. It was a huge labour, we were writing programmes in COBOL on an ICL 1906 machine, which was the university's mainframe, which reputedly although I think not truly took more than the domestic consumption of water in Oxford to keep cool, and would frequently time out while you were waiting for it to respond to the instructions that you gave it. So at that time it was pretty difficult. It got easier, and in the mid-1990s with a group of people at the IFS we managed to take this all the way back to 1961. That meant taking the computer tapes for the 1960s to a bit of Sainsbury's, because Sainsbury's was the only institution in the UK left with a mainframe on which you could read these computer tapes. It also meant sending a man called Steve Webb, who in the last government was the Lib Dem pensions minister, to the public records office at Kew, to dig out some more corrupted records which had been kept in paper. We managed to do it, and now, this really excites me and I'm sure it will excite you, we have a 53 year long consistent time series of this, and I'm going to show that to you. Here is 1961, on the same basis, and we're stepping through year by year. And you can see that through the 1960s there's nobody down here. There's a little bit of growth, so we're moving up a little bit, but not very dramatically. We're in about 1970s, 1970s, which were not a great period for the UK economy. There's still steady but not very exceptional change in the shape of this income distribution. Watch what starts happening in 1983, and keep an eye here, and here. We see an astonishingly rapid shift in the income distribution. The development of a fifth [?] tail, which we've not had before, just down there. Through the 90s, we go back to rather steady change again, and still an increase in average degrees of affluence, but nothing very dramatic, and the same is true through the 2000s. People are getting better off, but there isn't a radical change in the shape of this distribution. We're just coming to the end in 2013. Was this a big change? Well that's 1963, and this is 50 years later, and all of this is adjusted for inflation. So these are massive changes in the structure of the income distribution in this country. When did it happen? Well, not much happens between 1963 and 1973. So you can see that there is a little bit of thinning of the tail, and a small shift in median incomes, but no radical change. The same is true

between the next decade, between 63 [sic] and 83, not much happens. But then watch what happens in the following decade. For economists, this is an astonishingly big structural shift. We don't see changes like that very often, and it's only by looking at this kind of dis-aggregated, very detailed statistical work that you get a chance to see what's going on. In the following decade again, incremental changes, and the same this time round. So, a real, a really, very significant change in the structure of our UK income distribution, and one that you can only see by actually looking at the income distribution itself, by looking at the numbers, by doing the careful, and to some exciting, but to some slightly dull work of checking what's going on. So we had all this data, and we thought, what should we do first, and we looked at what was going on, and there was one feature of it that particularly interested us. And that was this really rather large number of people with zero incomes, more than half a million, and economists are very caring people, so we were very concerned about that. So we started investigating what was happening that could have led to this very significant increase in the number of people on low incomes, and we expected to find a worrying story. And we did find a worrying story, but it wasn't worrying in the way we thought it was, it was worrying in that it made us realise we needed to think more carefully about what we were measuring. So these are all measures of income. Of course, one of the characteristics of income is that it has to be measured over a particular period. It could be measured over lifetimes; this is actually income measured over a two week period. And what we discovered when we looked closely at the people on very low income is that they had really rather high consumption. So this is a chart that shows the relationship between consumption and measured income. We found that the people with zero incomes had consumption on average at levels as high as the people on beyond median incomes. Now, that's not suspicious, it's not that they're misbehaving. Well, some of them might be misbehaving if you take a particular Calvinistic view of the world. These are people who have had a good period, and so are not working for a while. So some of these are self-employed people, who've sold a lot of records, or whatever it happens to be, and are taking some time out from the paid labour market, but because they've accumulated assets and income in the past they are still able to consume at a reasonable level. Some of these are students, who are known for, ah, consuming in the legitimate expectation that their incomes will rise to meet their remarkable talents over the time to come. The reason I show you this is not that this is particularly important in and of itself, but because it's a very, very good example of how there are risks in all thought, not statistical thought, all thought, that we can end up focusing on something which isn't quite what it was that we really ought to be focusing on. Now, in the case of income, income is relatively straightforward to measure, so it is the right thing to measure, but you just need to be very careful when you interpret what is going on. Our first thought here was 'oh lots of people on low income, that's a sign of a terrible increase in poverty'; funnily enough it may be a sign of an increase in affluence, that there are more people who can afford to

have a period of very low income, because they have built up income elsewhere. Doesn't mean that poverty wasn't rising, actually it was rising at the time, but it reminded us to think carefully. It also helped us realise that when you're thinking about people on low incomes, what you're probably most interested in is people on persistently low incomes. So somebody who drops down to a low income for a single period, and in the period before and the period after has an acceptable income, then we might be much less concerned from a public policy perspective than people who we find in persistently low incomes. So this generated a whole stream of research activity that looked at volatility and persistence in low and high incomes. So, although it wasn't quite what we thought it was when we first saw it, it unlocked enormously valuable thought. And that's one of the characteristics of statistical work, that at its best, it encourages precision in thinking, because once you've measured something, you should at least know exactly what it is that you've measured, and if it doesn't behave as you'd expected it to when you interrogate it then that can help you think again. So this work, although it ended up with different answers to the answers we expected, I think turned out to be incredibly valuable, and it showed that this kind of aggregate data, which is the kind of data that often we look at, so this is just real national income, so the aggregation of all incomes in the UK, often will miss what's going on. So this is another very important chart. This shows what's happened to the real level of national income since 1948. So you can see that, well, you can't quite see but I'll tell you, the real level of national income is 5.53 times as high now as it was in 1948. We are unthinkably, in aggregate, better off, more than five times as well off as we were in 1948, and that's often forgotten. This is the simple product of compound interest. So, 1.0281 raised to the power 25 is two. So if something's growing at a little bit more than 2 and 3/4 percent a year it will double every 25 years, quadruple every 50 years, increase by a multiple of 8 every 75 years, and that's, by and large, what the economy has been doing. But the reason I wanted to show you this chart is that, I just showed you stuff that showed you that in the mid-1980s, from the mid-1980s on there was an astonishing shift in the income distribution. Well that's this period here. This period here looks like a period when the economy was growing, but it's not growing any more quickly than in this period here, this period here, or this period here. From this aggregate data which tells us what has happened to the overall level of income, we can't tell what's happening to the component parts of it, and that's rather important. The smart people in the audience will already have realised the answer to the 90/10 income ratio, because I showed you the 90 and the 10 first time round, and I'm sure that therefore the chart that I'm about to show you is completely unnecessary, but just for you to check your mental arithmetic, this is the 90/10 income ratio since 1961. So it's not a little bit less than four. If I go back to my notes, I'll see that 5 percent got the right answer, 20% got the answer adjacent to the correct answer, but between two thirds and three quarters got a completely wrong answer. Not slightly wrong, completely wrong, completely misunderstanding the shape of the

income distribution in the UK. Now, at one level that's quite funny, but another it's desperately serious, and, actually, this audience did slightly better than most audiences that I play this game with. You did better than the permanent secretaries for example, better than the No.10 policy unit in successive governments. And, while it's funny, it also makes me very cross, because it's really important to understand what's going on here, and it seems to me impossible to think coherently about public policy unless you know what the shape of the income distribution is like. Now you can see here that: this was the result of an enormous amount of work, and it's basically three straight lines, so it's a straight line from here to here, so through the '60s, '70s, and first half of the '80s, the 90/10 ratio is 3, it rises in the second half of the '80s to the early '90s to about 4 and a half, and since the early 1990s it has steadily declined, and is continuing to steadily decline, if I'd updated this, the next point would be down here. So this measure of the inequality of incomes shows flatness through the '60s and '70s, a very rapid increase in the second half of the 1980s, and a steady decline since then. Now, not only does the population at large, and the political classes, and the policy making classes not know what the 90/10 ratio is but they also don't know what's actually happened to income inequality. It's not true to say that income inequality has been rising, unless by 'has been' you've got a particular interpretation of the present tense which goes back to the second half of the 1980s. Now that's actually not ridiculous, because it's perfectly plausible that some of what we're seeing in the wider society is the lagged effect of that very, very significant increase in inequality. But if it is lagged effects that we're talking about, then we should be clear that it's lagged effects that we're talking about, not that we're trying to arrest a current increase in income inequality. Now it's worth saying that's it's just possible, so you can look not at this measure or at a Gini coefficient or at most standard measures, you could look just at the very top of the income distribution. So you could look at just the top 1%, and see a slightly different picture there, but even in the top 1%, for the last 10 or 15 years, there hasn't been a further increase in income inequality as far as we can measure. So, we need to know about this stuff, and it's not any longer very difficult to know that. What, just to remind you again, this is 83, this is 89, it was an extraordinary change, and I could speak for more than an hour about just why that happened, and honestly we don't fully understand why that happened, but it is something for us to look at. The consequence of that of course was that relative poverty increased. So this is poverty measured as the proportion of the population living on less than 60% of median income. Of course, if incomes on average rise without those at the bottom rising then you'll get a big increase in relative poverty. That's just what we saw. So in the second half of the 80s relative poverty rose very, very fast, and then has tended to fall away since then. But of course, funnily enough, relative poverty and absolute poverty will move in different ways. So the bottom line here is the line I just showed you, this red line is a measure of absolute poverty. So this is asking the question: what proportion of the population in each

year fell below 60% of the median income of the population in 2010, and you can see that in 1961 86% of the population had incomes below 60% of the median income in 2010. Strikingly as well, here is the period in the 1980s when relative poverty is increasing, but of course absolute poverty is declining. Because relative poverty was increasing because the median income was rising, and absolute poverty is falling, because the median income is rising. Now, neither of these is right, and my own view is that by and large it's relative poverty that we're more interested in, certainly in the rich developed nations, and so, this line which showed that dramatic increase is important, but it is also important to bear in mind what's happened to absolute measures of poverty, although of course that is largely simply reflecting massive increases in income.

So what does that mean that overall the government is doing? This is what incomes would look like in the UK if there were no government. Now that's a slightly ridiculous hypothesis, we can't, well some of us can imagine, it's tempting to imagine it sometimes; it's not a very serious counter-factual, because if there were no government then all kinds of other things would probably have to happen, but if you just take away the impact of government on incomes and on the provision of health and education, this is what it would look like. So the poorest 20% of the population having up to about£ 6-7000, the richest 20% having up to on average a little more than £80,000. But then government comes along and does a whole series of things, runs a health service, runs an education system, runs a social security system, and funds all of those with a tax system, and this is the consequence. So incomes at the bottom rise, those at the top fall. Overall this shows the government making us a bit worse off; that's a mistake. Or rather it's a failure perfectly to allocate all of these things. Look at it this way, this is the amount of redistribution. So, we take this money away from the top 20%, in particular the top 10% and 5%, a little bit away from this group, and allocate it to this. Now, I think there are all kinds of questions about how we do that, and whether we should do more or less of it, and more or less of it in particular ways, but again, it seems to be crucial that, if we're to have a discussion about that kind of thing, we have this kind of data available to us. Now of course, looking just at incomes like this is a very incomplete measure of what is going on, and one of the things that's happening at the moment which is changing the way we think about things quite a lot is that the probability of receiving an inheritance is rising quickly. As incomes rise, wealth accumulation across the whole of the income and wealth distribution is increasing, and so there is much more wealth to pass on. The probability of receiving an inheritance, we've got some people in the room this evening who are really wise and mature, and may be in this part, born in 1944. The probability of your receiving an inheritance is, well, unconditional on everything else about you, would be less than 30%, whereas youthful energetic people born in 1960, have got a nearly 60% probability of receiving an inheritance. And despicable long-haired, tattooed people with piercings, actually that's probably a bit early for piercing, although I have noticed that tattoos do seem to

be drifting up the age distribution – we don't want to see them – are up to 70%. And what that's a reminder to economists of is that the relatively simple world that we used to think of where you could think of individuals on their own in nuclear families, is less and less accurate a reflection of that world. This matters, but so does family structure. I can remember early in my career comparing myself with another alum of this college, Mike Devereux, who is now a professor in the economics department and runs the centre for business taxation here, and Mike was also working at the IFS, but he worked on companies and I worked on households. And I can remember thinking smugly 'how lucky to be working on households because, you know, they're straightforward. We know that the unit of analysis is a household, whereas if it's companies you have to think about subsidiaries, and international and so on. What's become increasingly clear over the last few decades is that many of those challenges that apply to companies apply to households too. We've all of us have members of our families, or friends, whose families are much more complicated, where there may be more than two parents involved in more than one child, and where, even the question that we assume there must be an answer to in the census, where does somebody live, becomes quite complicated, because if you have a child whose care is being shared absolutely equally between two locations, where are we going to say the child lives? Many of our assumptions about what is innate and inherent are being challenged, and we have to reflect that in what we're doing statistically, otherwise we will get into a terrible muddle.

Now, time is moving on. I was going to play a game with you, but I don't think I've got time to do it, so I'm not going to do it, but I'll tell you what the game would be. So I have here lots of dice, and the game I was going to play is a game that is all about, and the reason for playing this game is to explain how important it is to think carefully statistically. I was going to play a game about speed cameras. I'm a cyclist so I care about good driving behaviour. There are debates about speed cameras. The game I was going to play was this: imagine (we'll do a thought experiment), imagine a piece of road just outside where you live, and then we would use the dice to work out how many accidents occurred on that piece of road. I would give everybody one of the dice, and you would throw it twice, and that would be the number of accidents occurring on your stretch of road. We would then identify the accident black spots, and with the number of dice I've got which is about 70, I'd normally expect 5 or 6 people to score either 11 or 12, so to be accident black spots. Then I would take a piece of paper, I would draw a speed camera on it. I would give that camera to the people who scored 11 or 12. Holding the camera in one hand, and trying again to see if the camera was working, I would get people to throw the dice twice more, and see whether the camera works. Well of course the camera is extremely effective; it's very, very effective, normally you get a reduction of about 50% in the number of accidents. I'm sure you would have loved it, primary school children, they really adore it, they think it's great, think I'm Derren Brown, and that's the closest to Derren

Brown I'm ever going to get. Now, so what's going on there? A simple statistical thing which is called regression to the mean. So if you're doing something which is essentially random and the first time you do it you get a high number, the probability is that the next time you do it you get a low number, and so the speed cameras seem to work. So why play that game, because surely nobody would dream of assessing the effectiveness of safety cameras in that way, because that would be really stupid, it would reflect a failure to think statistically. The Department for Transport, early in the last decade, commissioned one of the better Russell Group universities to do some work on this, and what they did is they looked at places where there was an above average number of accidents, and they put speed cameras there, and then they measured to see what the effect on the number of accidents was. Well, they could have sacrificed chickens by the side of the road, and had almost all of the same effect, because, to the extent that accidents are caused by a heart attack, or a loud noise on the radio, or a clap of thunder, or a piece of mechanical failure on a car – that is, they're random, precisely this same effect would occur. Now, at that time, I was making the radio programme that Maggie referred to, and we made this point, to which people said 'certainly not'. Two years later, the Department for Transport published a revised version of the study, which said that 30% of the originally claimed reduction in deaths was the result of a long run downward circular trend. Correct. 60% of it was the result of regression to the mean. So they were left with 20% of the original number. Now, the then secretary of state was Alistair Darling, for whom I have a very high regard, and he made a good job of saying any reduction in deaths is to be desired, and of course it is, but if the mechanism by which you're reducing deaths is only one fifth as effective as you thought then you need to think about the other things you might do, like, changing lighting systems, putting fences outside primary school gates, and so on. So it is not acceptable to say that just because there's an effect in a particular direction it doesn't matter if you've completely miss-measured it. And that kind of thing, this was an extreme version, but that kind of thing does happen again, and again. And we must fight against it. And it doesn't require sophisticated mathematical skills to recognise there is a problem.

Now, I'm going to skip over a lot more of what I had thought about saying, I will talk about this just briefly. I'll talk about this and presentation. We hear a lot about the crisis of aging, and often we see these kinds of numbers, what's likely to happen to the number of older people in the UK over the next 20 years, and over 30 years from 2010, the number of 65-69 year olds is likely to go up by 40%, 75-79 by about the same, number of 80-84 by about 70%, number of people aged 85 and over is likely to double. I matriculated in 1978, that means that I hope to be in one of these two columns. I'll be this bar in the first half of the year and this bar in the second. I really hope that's the case because the alternative to being in those bars is being dead, and my preference would be to be in a histogram rather than be dead. Any time of the night or day I'd always rather be in the histogram

than six foot under. This isn't a crisis of aging, this is fantastic, this is the triumph of modern technology, that we live longer, you know, what's not to like? Now of course it would be possible that we were living longer but that the period of our lives that was associated with disability was growing, but even that doesn't appear to be the case. Disability-free life expectancy seems to be growing at the same pace as overall life expectancy. So why are we bothered about this? Well we're bothered about this because people are terrified of change. And one of the characteristics that's bad about these numbers is numbers are often used to show the change that's going to happen as though it's terrifying. The best antidote to that is the change that's already happened. So this is the chart that shows what's already happened to the number of people aged 85 and over, in the last century. So back here, in 1901 there were 61,000 people who were very old. By here, there were 1.447 million, if you believe the 2011 census, and the 2011 census was supervised by Sir Michael Scholar, so you better believe it. As far as I'm aware the UK has not been overrun by marauding Hell's Grannies over this period, British society and economy has not been destroyed by the consequences of all these wrinkly people, some of whom are here. And how lovely that is. All too often numbers are used to scare us by saying 'look, this change is going to happen'. Well, there's been a lot of change. I'm not going to show you modal ages of death, but I will show you this chart. This is what's happened to infant mortality rates in Sweden, and England and Wales. The Swedes, who are better at this than we are, have data going back to 1751, we only have the data going back to 1836. But in Sweden, in the latter part of the eighteenth century, more than 20% of children died in their first year. That rate was more than 200, it's now, well in Sweden it's probably less than four, in this country it's 4. It's even fallen in the lifetime of many of us here, there are people here, who when they were born, were relatively lucky, because 50 of every thousand children born then died. That number is now down to four. Astonishing change, and much the best way of capturing that change is with these kinds of numbers, indeed that's the only way of capturing the change. I'm not going to tell you about this, although maybe in questions.

Arithmetic. Arithmetic is really quite important. I could have picked any newspaper, this is the Daily Telegraph, this is headline in 2010: 'public pensions cost you £4000 a year'. Now that's quite a lot of money. They tried to do a good thing here. So there'd been a piece of work produced by the Office for Budget Responsibility that estimated that the cost of public sector pensions would be nine billion pounds. And they thought, well, nine billion, that's quite a difficult number, people don't find billions easy, don't know how many noughts there are, what is it anyway? Let's turn it into a number per person. But then they thought well no, we won't turn it into a number per person because there are more people than households, we want the number to be as big as possible, so we'll turn it into a number per household. So they said well, nine billion pounds, 25 million households. Divide one by the other. That's a really good strategy for explaining

numbers. Nine billion pounds divided by 25 million households. Now nine billion divided by 25 million is quite a tricky thing to do in your head, but let's make it easier. A billion is a thousand millions. So it's nine thousand million, divided by 25 million, so that's nine thousand divided by 25. You can already see where we're going. Nine thousand divided by 25 is the same as nine hundred divided 2.5. That doesn't look like 4000. So they had the answer out by an order of magnitude. And the reason that makes me cross is, if there were a misplaced possessive apostrophe, anywhere in the Telegraph, then one of the three people who read everything would notice it, and if they didn't notice it there would be hundreds of letters from people in Tunbridge Wells, saying, oh, you know, 'what's with the apostrophe?'. All of the people, this is the front page of the Daily Telegraph, they've got quite a few journalists, they've got quite a few people to check thing, nobody checked the calculation, and of course, if they'd done the calculation correctly they wouldn't have published the headline, because it wouldn't have been a headline, because it would have been a boring number. It wasn't boring, but it was garbage. And that kind of thing is happening again and again. And we are partly responsible, because with the whole of the rest of the population, we perform a kind of hypothetical surgery whenever a number comes in place, we take a hypothetical scalpel, and we slice it through the back of our neck, when we see a number coming, and we don't think about it, it's crazy. We have to start thinking about the numbers. I won't talk about admin data. I'll tell you a joke at the expense of statisticians. Censuses in the future will rely much more on administrative data, there's a lot of administrative data that is now accessible. Here is some work that some wonderful colleagues at the ONS did seeing whether you could by using the NHS population register and the Department of Work and Pensions customer information system which is essentially the National Insurance contribution record how close you could get to the numbers for the 2011 census. The more colour there is, the further away you are. So you see with those two systems you don't get very close and that's because in the case of neither system is there a very good incentive to clean the data, and you'll see that by and large, using these produces too many people, more people than you'd expect. If you add a third bit of information, student records, then you get very close. The joke at the expense of statisticians is that success here is measured by the bits of the country that are not brown, they're grey. Only statisticians would colour success grey. We are lovely, and we should be cherished, but we've still got some way to go.

Presentation. This is the last thing I'm going to talk about. This is really, really important. Here's a chart from a treasury document from a few years ago that shows the value of infrastructure expenditure in different areas, so communications, energy, flood, intellectual capital, transport, waste, and water. And one of the jobs that Michael had to do, that I had to do when I was chairman of the UK stats authority is that people would sometimes complain about statistics and their presentation. So who's worked out, raise your hand if you've realised

what's a bit funny about this diagram. Ok, so we've got a few people who've realised. What's funny about it is, this axis is all fine, but this axis seems a bit peculiar. So this says 1, 10, 100, 1000, 10,000, so this is a logarithmic scale. Now logarithmic scales have got lots of uses, but this is not one of those appropriate uses. Now let me tell you that it had been quite wet in the run up to this being published, and so there might have been an argument for this bar, 'flood' looking quite high. So that's a logarithmic scale. Would anybody like to guess what the relationship between transport and energy spending is on a normal scale? How much more do you think energy spending is than transport? Well, ask yourself that question, and then I'll show you the answer. So, transport spending is half of energy spending. I'm a pretty sad nerd, ok, I really am, I can recite the square root of two to quite a lot of decimal places, and I think that's quite cool, but I am not able to interpret that graph. I just can't look at it and know what the relationship of the true figures is. That's what the true figures are. And funnily enough, it shows that flood spending wasn't very big. So we wrote a letter about that. This kind of stuff really matters. It means you can't do the scalpel trick, you've got to concentrate, and it's not difficult to concentrate, but you have to retain that scepticism. Part of the reason for that is that the media, in this country and around the world, is not holding up a mirror. They're not reflecting what the world is actually like. This is the results from a brilliant piece of work done by a man called Roger Harrabin, who is now the BBC's environment correspondent, but for a while worked on health, and this looked at how many deaths you needed from different causes before you'd get a story in the newspapers. I think it was the Guardian and the Mail combined. If the media was holding up a mirror, all of these bars would be the same height. We need nearly 4,500 deaths from smoking before we get a story, 846 from alcohol, 2,500 from obesity, 375 form measles, one and a half deaths from human variant CJD, 22.5 from HIV, and 1375 from mental health. And again, it's quite funny, but it isn't funny, because these are not slightly skewed, they are a complete misrepresentation of reality, and yet that's where we get most of our understanding and information from, that, and anecdote. The country, the world in which we live is too big for anecdote to do it anymore. The people that you meet in Oxford, or London, or Manchester, or wherever you live, they will not give you enough information for you to understand what the world looks like. If you want to understand what the world looks like, you've got to go out and measure it. Measuring it is difficult, it causes all kinds of difficulties, and we often get it wrong, but it is better than the alternative, because the alternative to measuring it is that we are driving the car with the windscreen blacked out, and, more often than not, that's what the world feels like. These forms of measurement, they are beautiful, they are powerful, statistical thinking is a wonderful aid to precision in thinking, and if we don't concentrate and take it seriously, then we'll make many more mistakes than we possibly ought to, and John Kay wouldn't be a happy man, and that's much the most important thing on a night like tonight.