Chapter 2: Page Modeling System

Chris Hope

I'm going to talk about PAGE2002, which is the integrated assessment model used in The Stern Review, and to talk about what it can tell us about the social cost of carbon, which, as you know, is the impact of one more ton of carbon being put into the atmosphere now in the form of CO2. At the end I'll talk a little bit about some of the other things that the PAGE2002 model can do for us as we try and decide what to do about this serious problem.

I'm going to start with the structure of the model itself so you can see what's involved in it, talk a bit about the values that come out for the social cost of carbon in The Stern Review, and explain where those numbers come from. One of the advantages of the model is that it is transparent, and I want to make those numbers transparent to you so that we can understand where they've come from.

Then I'll lead into talking about what happens to the social cost of carbon, the impacts of a ton of carbon emitted as carbon dioxide, if we make some other assumptions based on some of the comments that have been made. This leads us into the discussion that we'll be having this afternoon. Then I'll talk for a few minutes about what else the PAGE2002 model can tell us.

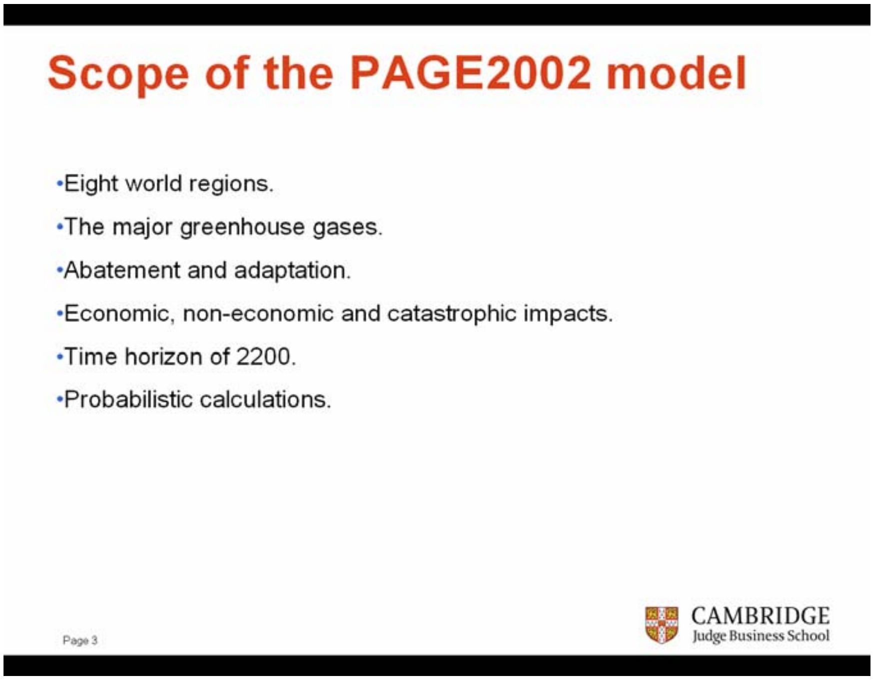

So, very simply, the PAGE 2002 model is an integrated assessment model. It divides the world up into eight regions. It looks at all the major greenhouse gases, not just carbon dioxide. It allows you to make decisions about abatement, how you're going to cut back your emissions of gases, and adaptation, how you're going to cope with any impacts that might be caused.

It looks at economic, non-economic, and catastrophic impacts, the kinds of things that we should be very worried about if we get temperature rises above four, five, six degrees Centigrade maybe. It has a time horizon out to 2200 because this is a hugely difficult and long-term problem.

And, most importantly, Sir Nicholas emphasized uncertainty all the way through his talk. All of the calculations in the model are done with a recognition of that uncertainty. So, they're done probabilistically. You don't just put in single values for parameters and get out single answers. You put in ranges for parameters and get out ranges for answers.

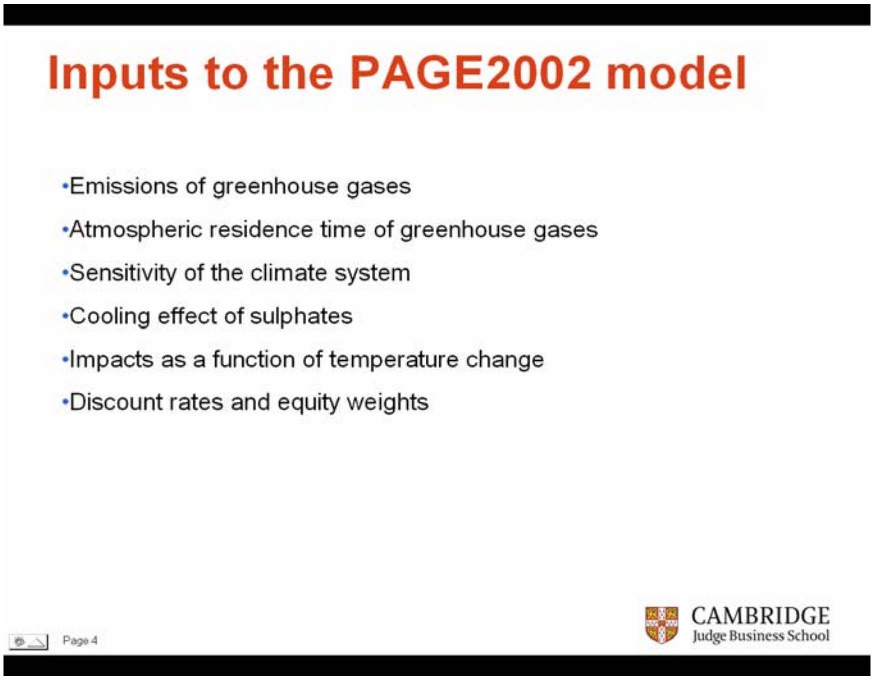

What kinds of things do you need to put into the model to be able to do these calculations? You need to have some projections of what the emissions of greenhouse gases are likely to be. You need to make some assumptions about the atmospheric residence time of the greenhouse gases. How long are they going to be up there? How sensitive is the climate system? Nick has shown a slide on the assumptions that are in there about that. You need to look at sulfates and the cooling effects of those.

So, this is sort of scientific information that needs to go into the model to be able to do the calculation, but also you need to look at impacts as a percent of GDP and how those vary as a function of the temperature changes, and you need to look at the discount rates and equity weights, the kind of economic parameters that Nick was talking about. This is what makes it an integrated assessment model. It puts in values from the scientific studies, values from the economic studies, brings them together, and allows us to say what kind of implications they have for the policies we might want to follow.

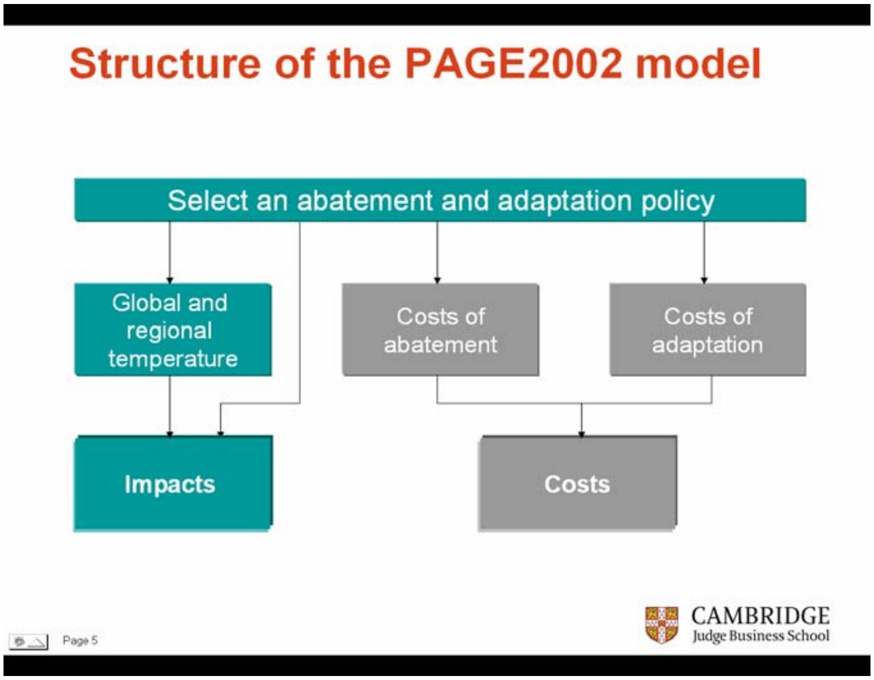

The structure of the model is that to run the model, you need to make some decisions about what your abatement and adaptation policy is going to be in all the different world regions, how much of each gas are you going to allow to be emitted, what kind of adaptation are you going to have. You can then calculate through on the left-hand side, shown in blue, which works out the global and regional temperatures that result from the emissions of the different gases. And if we look at the amount of adaptation, that can then, in combination with the temperatures, tell us what the impacts are likely to be. I won't talk very much about the things on the right-hand side of the model, which allows us also to say what are the costs of adaptation, what are the costs of abatement, and get the costs of taking action as well, but I will refer to that at the end.

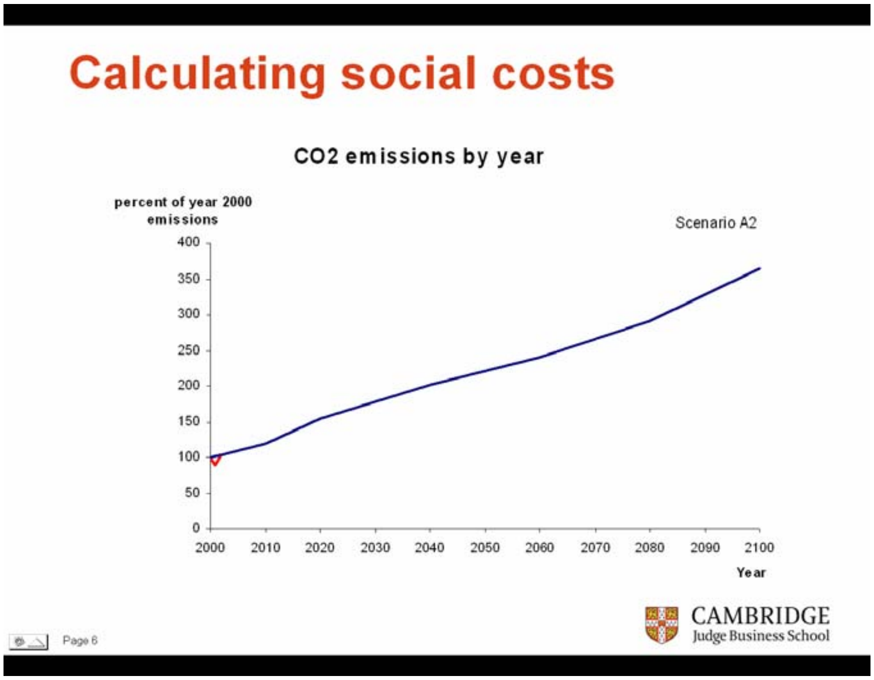

So, what do you actually do if you want to work out the social costs of a greenhouse gas, in this case, carbon dioxide? Here's a slide showing what the emissions might be over time. Under scenario A2, one of the IPCC scenarios, you can use the model to calculate in each region in each time period what the impacts are likely to be in a probabilistic sense for those emissions, and then you can aggregate them up over the regions and aggregate them back through time and aggregate them across the different outcomes, as Nick was talking about, to give an estimate of what the total impacts are likely to be if you have emissions along that blue line there.

And then you can do exactly the same thing with emissions just the same as along the blue line except for that tiny red triangle taken off at the beginning, and you can work through again all the impacts that there will be in the different regions over time over the different outcomes with that scenario exactly the same, except for that little red triangle there, and then you can take away one from the other.

And as that little red triangle actually represents one billion tons of carbon, if you then divide the difference between the two scenarios by one billion, that will then tell you how much extra impact is being caused by one ton of carbon being emitted today. And that's the way these kinds of calculations are done to work out the social costs of the marginal impacts.

So, what kinds of figures do we get from The Stern Review? The number for the social cost of carbon that is most prominent in the review, a figure of $85 per ton of CO2, translates to $312 per ton of carbon, which is the more usual unit that people measure it in. As Nick said, this is being calculated using the PAGE 2002 model with sets of assumptions in there and transforming the answers using balanced growth equivalents and so on afterwards in order to be able to produce this expected value number.

Now, the first thing we can do with the model is say, well, okay, let's make all the assumptions exactly the same as in The Stern Review and let's rather than just produce one number, the central number, let's see what kind of range of numbers comes out of the model. And this is what's shown on this slide.

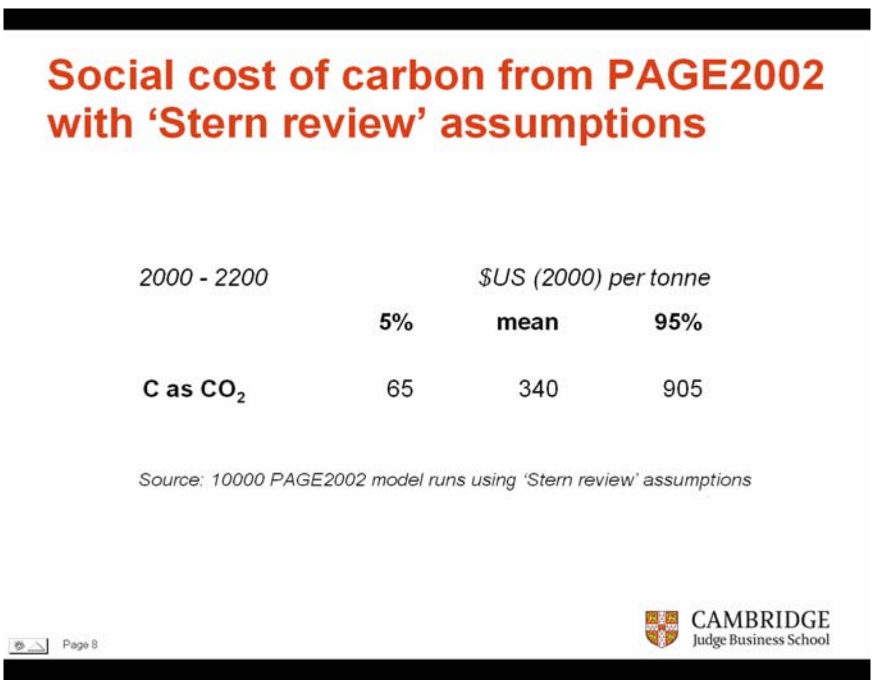

I've put 'Stern Review' here because this isn't the answer using the clever balanced growth equivalent mechanism that Nick and his team have been able to do. This is the answer just taking the numbers in terms of the effect on GDP, transformed using equity weights, and you can see that we can end up reproducing the answers from the Stern Review pretty closely. The mean value that comes out of the model when you run it like that is $340, just a bit higher than the central number produced in the Stern Review, and the five to 95% range is $65 to $900 per ton of carbon.

So, immediately you see just how uncertain we are about this. Even using all the parameters pretty much exactly the same as in the Stern Review, we end up with a 90 %confidence interval for the social cost of carbon that is an order of magnitude, or a bit higher. One of the advantages of having a probabilistic model like PAGE is that you can immediately get out of it some information about which of all the uncertain inputs there are in the model have the biggest effects on the answer that you're getting out. In this case, the answer for the social cost of carbon. And we can see from the next slide that it's a mixture of the scientific information and the economic information that goes in there.

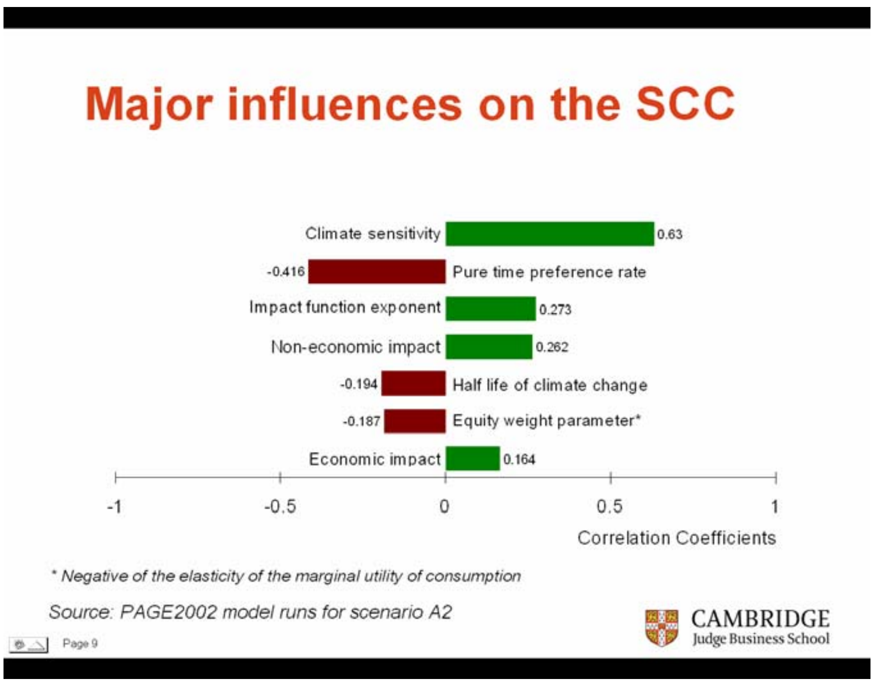

The most important parameter, the one with the longest line here, is the climate sensitivity. Our uncertainty about that has the biggest impact on the answer we get for the social cost of carbon. The next most important is the pure time preference rate that you assume and so on down. The third one is the impact function exponent that Nick spent some time talking about as being an important parameter, and we can see from this slide that it is one of the top parameters here.

So, this slide tells me which parameters I need to tell you about to get you to understand why the answers in the Stern Review come out as they did.

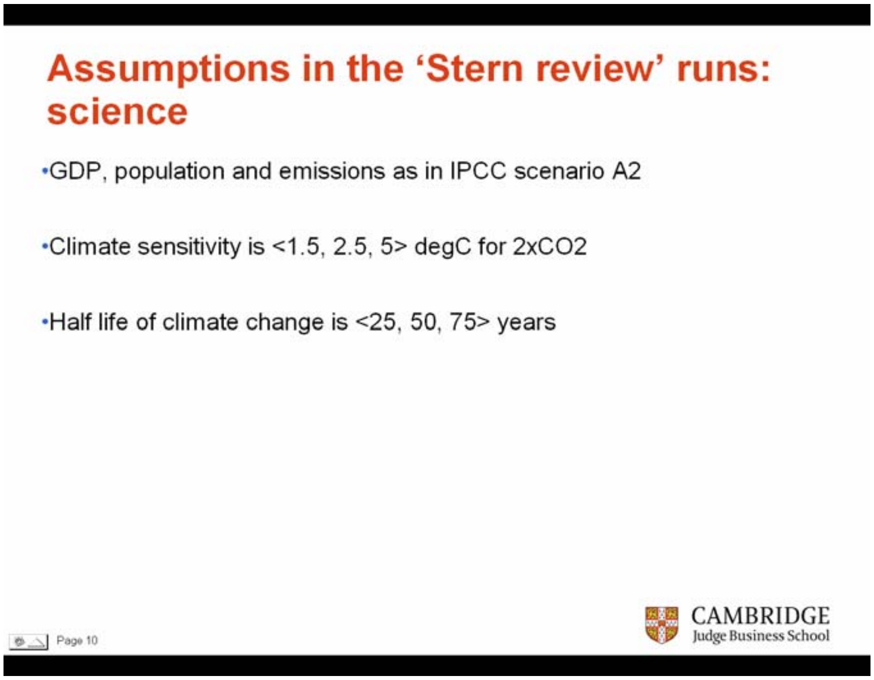

The assumptions in the Stern Review runs, which ended up with that range of social cost of carbon that I showed a moment ago, are these: The GDP, the population, and the emissions are as in the IPCC scenario A2. Climate sensitivity has a range somewhere between 1.5 and five degrees Centigrade for a doubling of CO_2 concentrations, with a most likely value of 2.5 degrees. The little triangles around the numbers show this is a triangular distribution with the bottom point being the minimum, the top one being the maximum, and the middle number being the most likely input.

So, as Nick said, this is a very standard assumption about climate sensitivity, standard assumptions about the half-life of climate change, how long it takes on average for the earth to respond to any change in forcing of between 25 and 75 years.

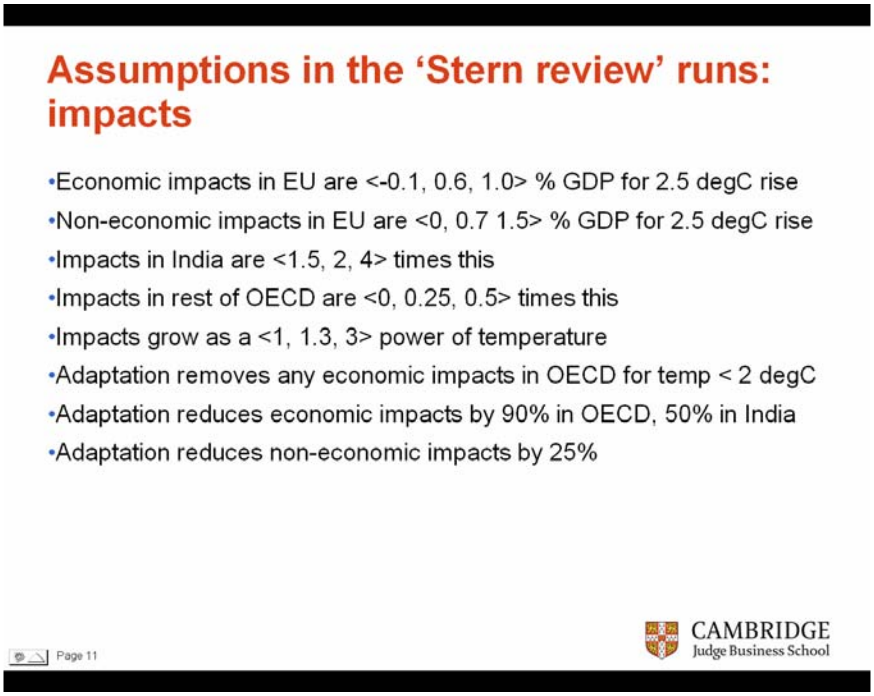

The impact inputs. The economic impacts in the European Union in the model in The Stern Review runs are assumed to be somewhere between minus 0.1% to plus 1% of GDP for a two-and-a-half degree Centigrade rise. In other words, the model does allow a small possibility that global warming could actually be a good thing, could lead to benefits rather than costs in regions like the European Union.

However, the non-economic impacts, the effects on health, the effects of inundation of wetlands and so on that don't get directly into GDP, are assumed to be somewhere between zero and 1.5% of GDP for a two-and-a-half degree Centigrade rise.

Now, the impacts in India or other poor countries are greater than this. There's a multiplicative factor that's maybe twice as high in terms of percent of GDP lost. On the other hand, the impacts in the rest of the OECD are smaller than in the EU, because in Europe the coastlines are long and coastlines are quite vulnerable to things like sea level rise. Other assumptions are made about how impacts vary as the power of temperature. Well, here they are. The power function is somewhere between linear and cubic in temperature, with the most likely just slightly more than linear in temperature. As temperature goes up, impacts go up to the 1.3 power of temperature.

There are some fairly aggressive assumptions about adaptation in the model. Adaptation is assumed to remove any economic impacts in the OECD for up to a two degrees centigrade temperature rise. It's perfectly adapted to temperatures up to that level. If we go beyond that, it's assumed that adaptation can reduce impacts 90% compared to what they're stated as in the first line of the slide. Adaptation is not so effective in poor countries, but it still reduces economic impacts by 50%. The non-economic impacts can only be reduced by 25% because it is not to see how we might be able to prevent things like the inundation of wetland.

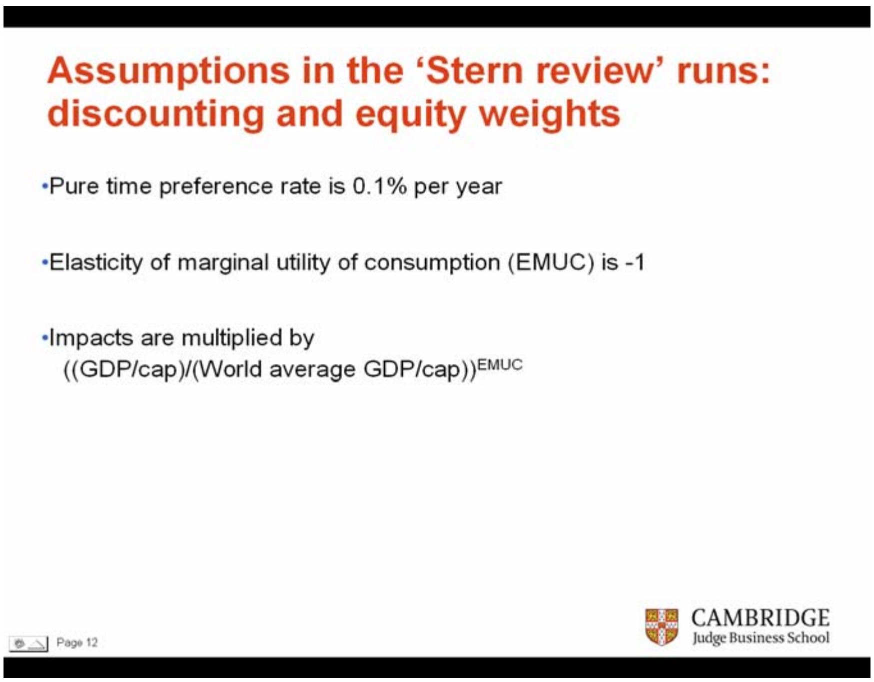

The discounting assumptions are the ones that have attracted most comment and Nick has talked about the pure time reference rate in these runs of 0.1% a year. The elasticity of marginal utility of consumption is the same thing as Nick's eta parameter, set at minus one and the equity weights are applied by multiplying the impacts in each region by that equation at the bottom. The effect of that equation is that a dollar's worth of impacts in a poor country is valued at more than a dollar's worth of impact in a rich country.

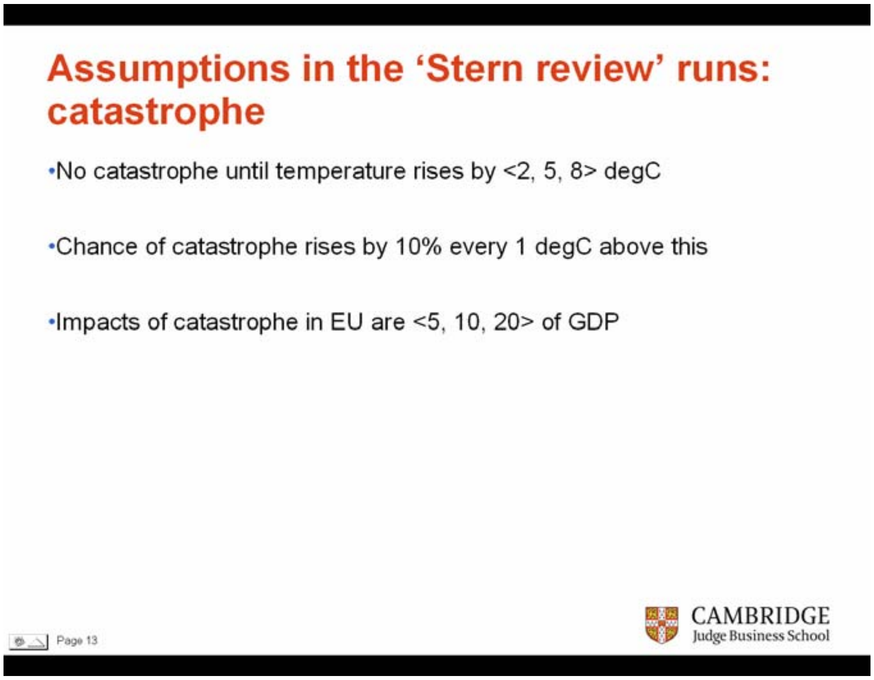

Finally I want to show you what kind of assumptions are made about the possibility of catastrophe. Catastrophic impacts are things like the melting of the West Antarctic ice sheet, and are what many people are most worried about in the long term. It's assumed there's no chance of this kind of catastrophe if the temperature stays below about five degrees with a range of two to eight degrees above pre-industrial levels. If you do get above that threshold, the chance of a catastrophe rises by ten %for every degree centigrade. And if we do get a catastrophe, that's quite severe; in the EU you lose between 5 and 20% of your GDP, with a most likely loss of ten percent.

So they're the main inputs to the model. If you use these inputs, and I don't think they're outrageous or outside the range of numbers that come from the literature, you will end up with a social cost of carbon of about $300 dollars per ton with a range of maybe $60 to $900 a ton. It's just a consequence of the inputs. Of course, if you don't think those inputs are right, then you will end up with a different social cost of carbon and therefore different implications for what you might do. What I'm going to do is just run through a few changes that you might make to the input numbers and see what effect they have on the social cost of carbon.

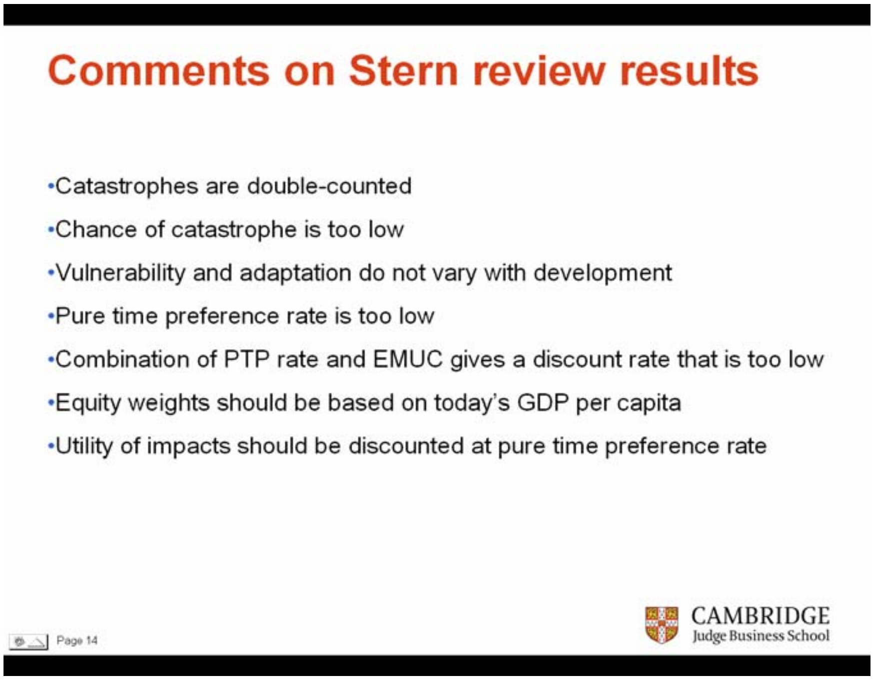

So, the kinds of comments that there have been—and this leads into this afternoon's discussion—on the results are, well, maybe this estimate of catastrophe that we've got in the model is not quite right. Maybe it should be either lower or higher, and maybe the assumptions we've made about pure time preference rate and the elasticity of the marginal utility of consumption are wrong; they end up giving you a discount rate that's lower than we observe in the market place.

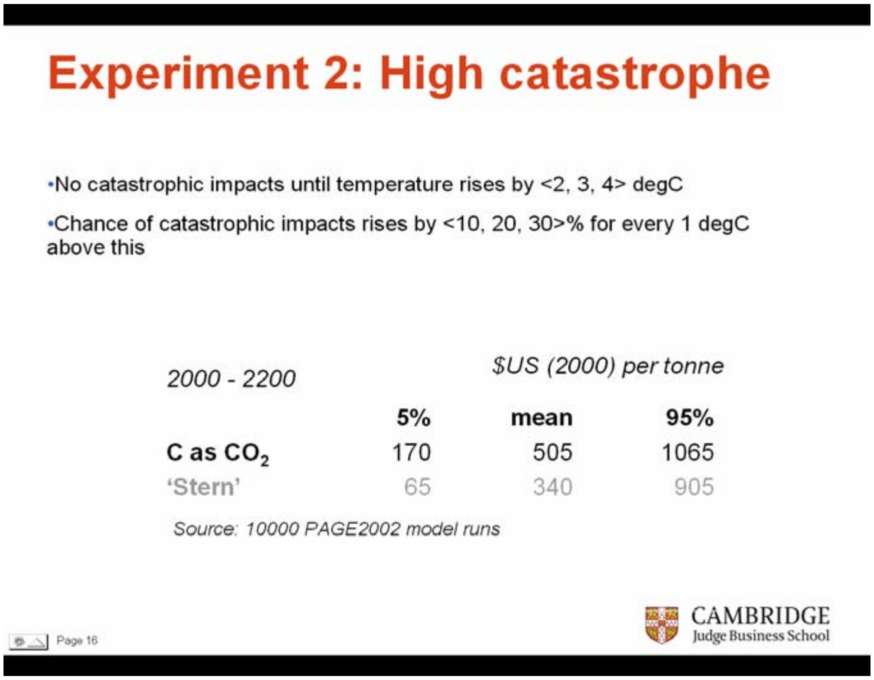

The assumptions that we made for catastrophe in the Stern Review runs end up giving you maybe a 2 ½ % chance there's a catastrophe by 2100 and a 25 % chance that there's a catastrophe by 2200. Some people have said, actually, it's more likely than that.

Here are some parameters which try and take that into account. If instead of saying we can't have any catastrophe until the temperature rises by five degrees, we say, well, maybe we could start having a possibility of a catastrophe if it goes above three degrees with a range of two to four, and if the chance of a catastrophic impact rises not by 10 % for every one degree centigrade above this but by 20%, then we end up with values for the social cost of carbon about 50 % higher. $500 is the most likely value rather than $340.

Even though—even with these new values for a higher possibility of a catastrophe, there's still only about a 25% chance that you'll have one by 2100 and about an 80 %chance you'll have one by 2200. So, you're talking about things that are most likely to occur in the 22nd century, over a hundred years from now, and yet they still have quite a big effect on the social cost of carbon today.

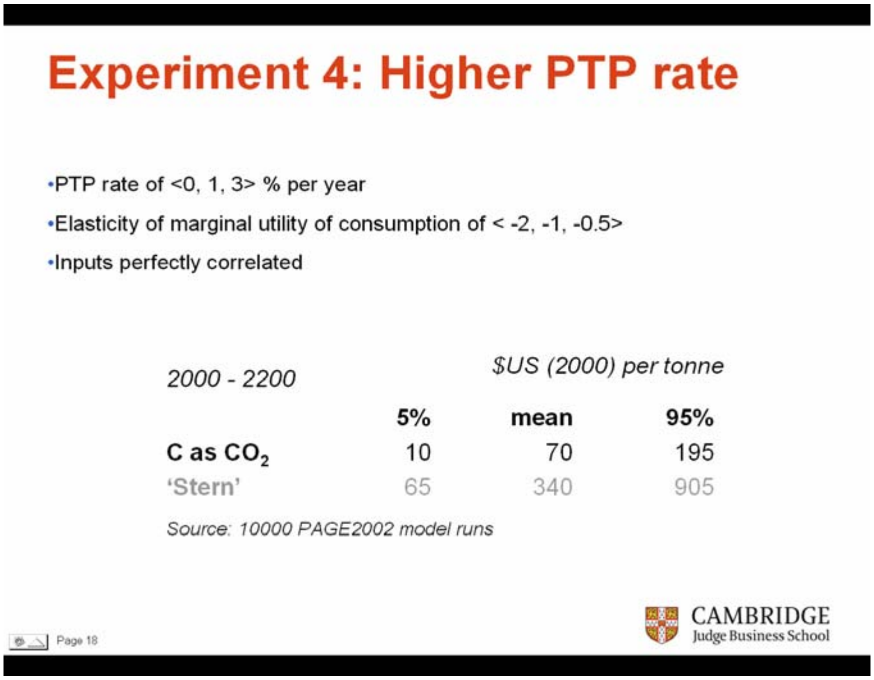

What happens if you take a different ethical position from Nick and his team? What happens if you assume that the pure time preference rate is a bit higher, maybe somewhere in the range of zero to 3% a year, most likely one percent? And what happens if you assume the elasticity of marginal utility of consumption again has a range going up to minus two, the kind of number that Nick was talking about and maybe down to minus 0.5? And what happens if you assume those inputs are perfectly correlated so that you end up with the kinds of consumption discount rates of the order of 4 % a year that we tend to see in the economy as a whole? Well, the answer then is if you use exactly the same scheme for doing equity weighting and discounting, you end up with a social cost of carbon that's about four-fifths lower. The mean value goes down to about $70 per ton of carbon and the range from ten to about $200.

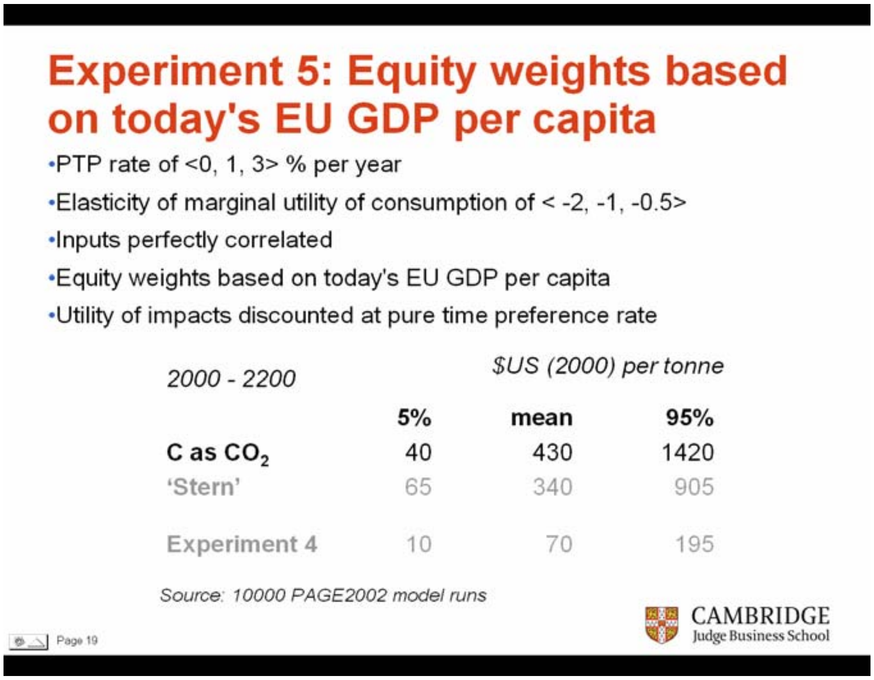

But that's not really the end of the story because if you were wanting to use this kind of experiment about higher pure time preference rates and range for the elasticity of marginal utility of consumption, you would probably want to go the whole hog and do the kinds of different runs that a lot of economists have suggested we should do, where you base equity rates not on the world average GDP per capita in a particular year, but you base them on today's EU GDP per capita, and you then discount your utility impacts not at the consumption discount rate but just at the pure time preference rate.

And if you do that, and keep everything else exactly the same as experiment 4 which I just showed you, the mean value not only comes back up from $70 per ton of carbon back up to $300, but it actually goes higher, to a mean value of about $430 per ton of carbon.

So, there's clearly a lot to be discussed—and I know we're going to discuss this more this afternoon, and that will be good because one of the advantages of having a model like this around is that when you get different views on what should be put in for different parameters, the model's there, you're able to run them through and see what implications they have for the social cost of carbon and, therefore, the policy that you should take.

What I'm going to do now is just run through in the last three or four minutes some other things that you can get out of this kind of integrated assessment model. These are results that come from runs done before the Stern Review. So, typically, with higher pure time preference rates than the 0.1 %per year that Nick's team has used and, therefore, the social cost of carbon that you get from those kinds of runs might be down around $40 as a mean value.

Social costs of other gases

2000 - 2200	*$US (2000) per tonne*		
	5%	**mean**	**95%**
C as CO_2	10	43	130
Methane	40	205	500
SF6	70 000	520 000	1500 000

Source: PAGE2002 model runs for IPCC Scenario A2

Page 20

CAMBRIDGE Judge Business School

Because you're putting all the other greenhouse gases in the model as well, you can work out what's the social cost of the other greenhouse gases, too. You can work out how much emphasis you should be placing on reducing methane, reducing sulfurhexafluoride, or the other gases compared to reducing carbon. If we're going to solve this problem, we're not going to just do it by tackling carbon dioxide alone. We're going to have to tackle the other greenhouse gases as well.

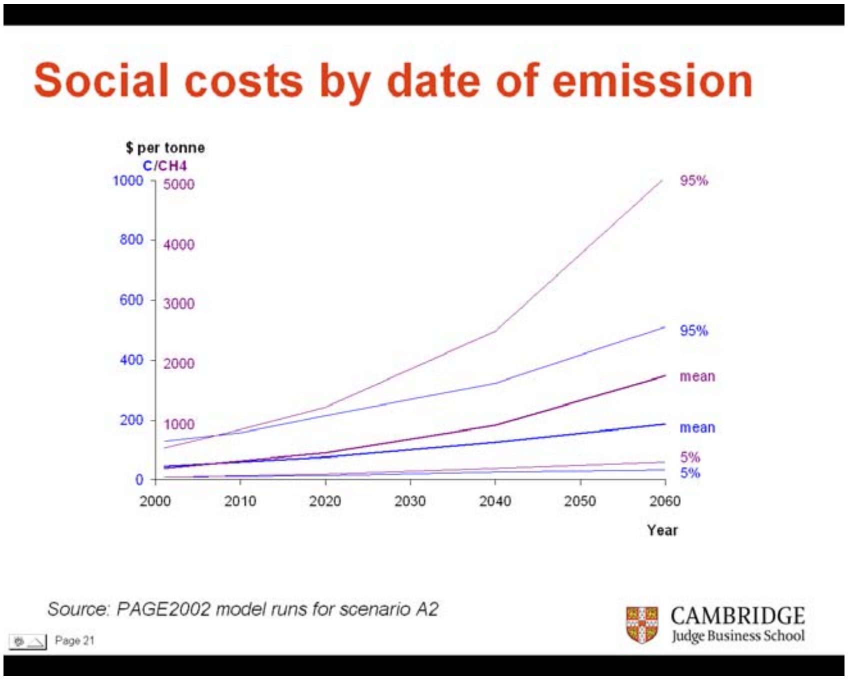

The model can also tell you if you take that little red triangle of emissions away—not in the base year today, but at some point like 2010 or 2020 or 2040—what happens to the social cost of carbon over time and the social cost of other greenhouse gases over time. And here we can see in blue the values for what happens to the social cost of carbon over time and in purple what happens to the social cost of methane over time, and the scales are calibrated such that they look equal in the base year today.

And you can see that what happens is that they both increase over time. They increase because as we get closer to the time when you would expect the most severe impacts of climate change to occur, then the extra impact that you'll get from putting another ton into the atmosphere gets higher. The mean value for the social cost of carbon goes up at about 2 ½ % a year. The mean value for the social cost of methane goes up more than that at over 3 % a year, and that's to do with the different lifetimes of the different gases in the atmosphere. I can explain it in the discussion if people want to know more about it. So, that's another thing you can do.

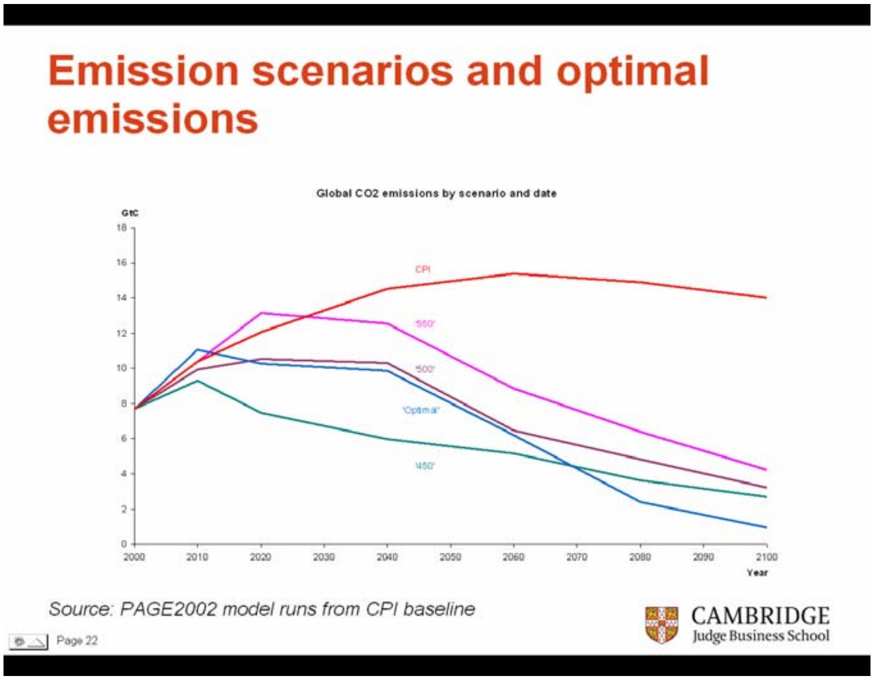

We need to think not just what should be used for the social cost today, which might set us a carbon price or a methane price today, but how is that going to have to change over the future as we carry on with this policy that's going to have to be there for decades if we're going to make any impact on this. If you remember the slide showing the structure of PAGE that I had right at the beginning and say, okay, let's look at the costs of doing something about this problem, the cost of adaptation and the cost of abatement, as well as the impacts, you can add those two things together. You can add together the total impacts that are left from climate change and the cost of doing something about the problem, and you can then begin to say, okay, what kind of policies can we put in place that will try and minimize those costs, again aggregated over space and time and all the different outcomes in a probabilistic way, and you end up perhaps with some lines like that blue optimal line on the graph.

That's the sort of set of policies that might end up minimizing the sum of the action cost, the cost of doing something about the problem, and the impact cost, the cost of climate change as a whole. And you see that even with much lower values for social cost of carbon than we get in the Stern Review, those kind of optimal policies are likely to end up leading you down to quite low emissions as we go further into the future.

And if you have that, if you can do something clever with genetic-algorithm-type methods which allow you to work out an optimal path of emissions even under uncertainty, then you can do something which is really quite valuable for the research community as well and say, okay, that's our optimal path if we don't know any more about the problem. If we've got this huge uncertainty that we have at the moment, it's the blue dotted line there.

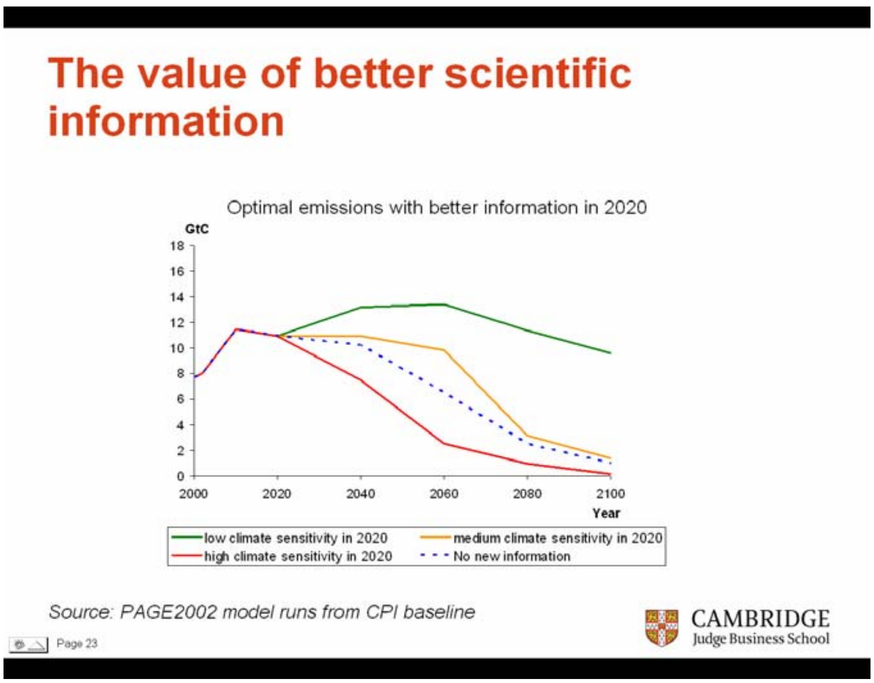

What happens if we get some better information in, say, 2010, or in 2020? Well, this slide shows what happens if we get better information about the climate sensitivity, this very important scientific parameter, not that we get perfect information, we know exactly what it is in 2020, but it's just a bit better. We reduce the range by about half compared to what it is today.

Obviously, if you get information that the climate sensitivity is towards the low end of the range, then the optimal emissions will be higher, the green line there. If you get information that the climate sensitivity is towards the high end of the range, the red line, your optimal emissions will end up being lower. And by looking at the total cost of impacts and costs of taking action once we've got this better information, you can work out what the value of this information is to you.

It actually turns out that the value of this information is enormously high. If we get information that halves the range of our climate sensitivity parameter in 2010, that's worth about 400 billion dollars to us. If we get that information in 2020, it's worth about 300 billion dollars to us, not quite as much because we aren't able to take action quite so early.

But that gives you a huge impetus for thinking about what kind of research you should be doing and how much you should be spending on these kinds of scientific and economic studies that we need to do. It also, of course, assumes that people take notice of them and will adjust their emissions optimally as a result of them, but, you know, we all have to make some kind of assumption there.

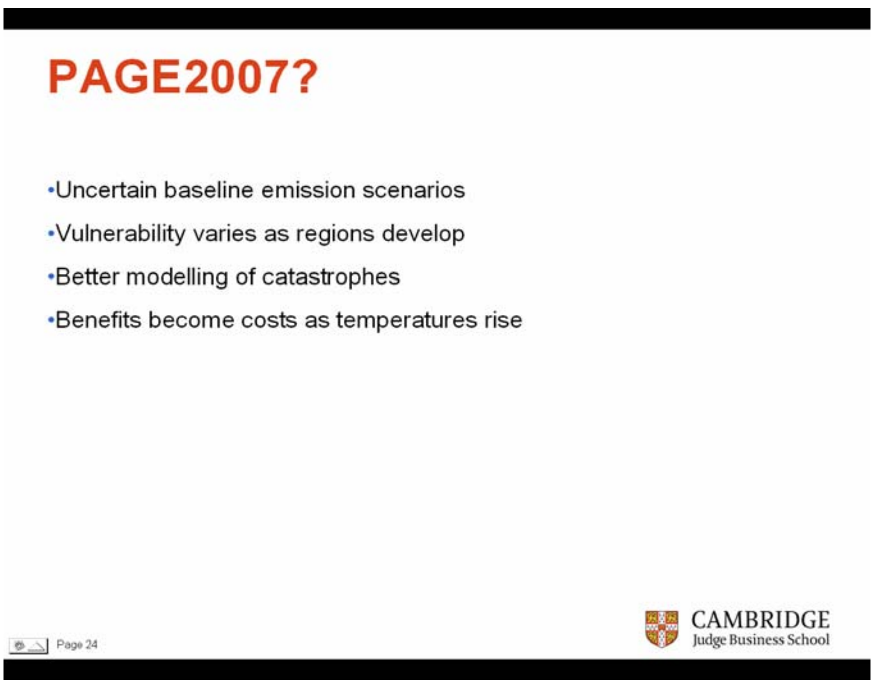

So, that's it really. What I've tried to do is be as transparent and open as I can about what goes into this model that's at the heart of the review, calculating these kind of numbers, show that the numbers that are in there are not outrageous, they're fairly standard, and what happens if you change some of them. Of course, the model isn't perfect. It carries on being developed. As the new IPCC results come out this year, we'll want to have new versions of it and keep it going.

If anybody has any questions about what I've run through here or would like to think about using the model themselves to run through their own assumptions that they might have, which might be different from the ones I've shown here, well, that's the great advantage of having the model there. It means we can do that kind of run and see what different assumptions do to the policies we would want to adopt, and I'd be very happy to work with anybody who might want to do that.