Chapter 1: Findings of the Stern Review on the Economics of Climate Change

Sir Nicholas Stern

In June 2005 the Chancellor of the Exchequer, Gordon Brown, asked me to conduct a review of the economics of climate change. This was not to be like the reviews that I and many of you have done of academic literature, in which you try to put a whole literature in perspective, give it a structure, and try to be exhaustive and fair to everybody who's contributed to that literature.

Our task was to produce a different kind of review—one that would help people involved in policymaking come to a view on what would be a sensible way forward, given what we now know. So, it's a review that speaks to policy, and is mostly about policy.

To make suggestions about policy, you have to come to a view on what kinds of things should be done. In the case of climate change, how strong should the action be? How urgent is it? Those questions are the first steps, and then you look to what kind of policy tools can steer you in a sensible direction. Thus the review has two halves: the first is about what kinds of actions are necessary in the face of climate change, on what scale, and when, and the second half looks at the details of policy design. I'm guessing that the participants here today are going to take the modeling side of the Review quite seriously, so I want to talk about that, too. But as all modelers should know, formal modeling is only one part of the argument and one part of the contribution of economics. Thus I hope that in the discussion we can get into the policy issues more deeply.

Why Bother About Climate Change?

First it is worth thinking about why one should bother about climate change at all. I have to say that I came to this question in July 2005 with a very open mind. I knew what the greenhouse gas effect was, but I hadn't thought through, or been closely involved in discussions of, what might be sensible policies to approach it.

Three reasons are advanced for why one can relax and forget about climate change. I think they are fundamentally misguided, but it's important to recognize them. One is that the science is all wrong, that global warming is "the biggest hoax ever perpetrated on mankind," similar to how one US senator has described it. That is for the scientists, but the weight of evidence is now such that most people would see this point of view as simply absurd. The February 2007 Fourth Assessment Report of the Intergovernmental Panel on Climate Change (IPCC) sets out the evidence in a very convincing and clear way.

The second reason for denying the problem of climate change is to say that human beings are fantastically adaptable; they can cope with anything that comes their way. Human beings are, of course, adaptable, but if we go on with business as usual, the risks we now run, for some time next century, are for global warming by five, six degrees Centigrade and above, relative to pre-industrial times. Such a rise in temperature would be earth-transforming. Five degrees Centigrade is the difference between now and the last ice age 10 or 12 million years ago. The kinds of places we could live and how we could live our lives would be radically transformed in ways that are very hard to understand. For one thing, these kinds of transformations involve very big movements of population, and we know from the experience of the last century or so that the world hasn't got any better at handling big movements of population. So, I think the idea that we can adapt to anything that comes our way is reckless, relative to the kinds of things that could happen.

And the third reason for denial is that "Whatever the effects are, they're way off in the future, and I'm not particularly bothered about the future, so I'm not going to do anything." That attitude involves what some economists would call pure time discounting, and I'm convinced that such time discounting at a heavy rate would be viewed by most people as unethical. It involves discrimination between individuals by date of birth. The ethics of climate change is a discussion that we ought to have. Whilst markets do have some information, it is not easy to use without very strong and fairly implausible assumptions. And there's no way we can simply read off the relevant ethics from the behavior of markets: we cannot see a collective expression in the markets of what, acting together, we should do for 100, 150 years' time. This is an area in which reasonable people can differ, but in which reasonable people are obligated to have a serious discussion.

The economics of climate change is fundamentally about the economics of risk. And if you act on climate change and invest in bringing forward new technologies, and it turns out to be the biggest hoax ever perpetrated on mankind, you will still have acquired a lot of new technologies that are probably quite useful. If, on the other hand, you assume it is the biggest hoax ever perpetrated on mankind and you do nothing, then you may fairly quickly end up with a lot of irreversible and severe damage from which it will be very hard to extricate yourself.

The Economics of Climate Change

Let me now move directly into the economics. If you emit greenhouse gases, you cause damage to other people, you influence their ability to produce and consume, and that's what economists mean by an externality. London dealt with the familiar externality of traffic congestion by the price mechanism: it introduced a congestion charge for vehicles entering the middle of the city, though not until the average vehicle speed had slowed to walking pace. Delhi moved to green (actually CNG) auto rickshaws and buses by government fiat, after people had been dying very heavily of bronchial and other diseases and the exhaust-gas problem had been elevated into a question of human rights: the Supreme Court in India intervened.

But climate change is an externality in a very different form from what we are used to. It is global in its origins, and global in its effects. Greenhouse gases emitted in Australia and New Haven and London all have the same effect on the atmosphere. The impacts occur throughout the world, though of course, they differ in different parts of the world. Climate change also occurs over the very long term, and this means that we can be in a crisis without actually seeing the direct effects immediately—there is no equivalent of directly experiencing the cars slowing to walking pace in London. So, politically, it can be difficult to get action. Climate change is also very uncertain. We don't know exactly how much emissions will arise from different types and levels of economic activity or how much a given concentration of greenhouse gases raises temperature. We don't know exactly what effects different rises in temperature have on weather patterns. And we don't know exactly what effect weather patterns of different kinds are going to have on production and consumption. There's lots of uncertainty right through this chain, and it has to be taken into account directly. And, of course, the effects of climate change are potentially very large and irreversible.

Now, what are these effects? Mostly they happen through water: through storms, droughts, sea-level rises, and floods; for example floods of the kind we saw in Mozambique in 2000, which took many thousands of lives and knocked 15 % off that country's GDP. Some effects happen directly through heat, like the heat stress we saw in Europe in 2003, in a summer that will probably be normal by 2050.

Such events occur in a stochastic way but can be very big and very disruptive. What we're seeing now, of course, is only on the basis of 0.7 degrees Centigrade of warming, relative to pre-industrial times, rather than the kinds of effects that the world is going to experience as the temperature rises further. Even if we start acting strongly and urgently now, we shall probably get three or four times this amount of temperature rise. Of course, if the rise goes up to five or six degrees Centigrade, we shall see something that is quite hard to imagine. Figure 1 illustrates, with the arrows reflecting growing risk and impact, how much worse the effects will become if we go up into these higher temperatures.

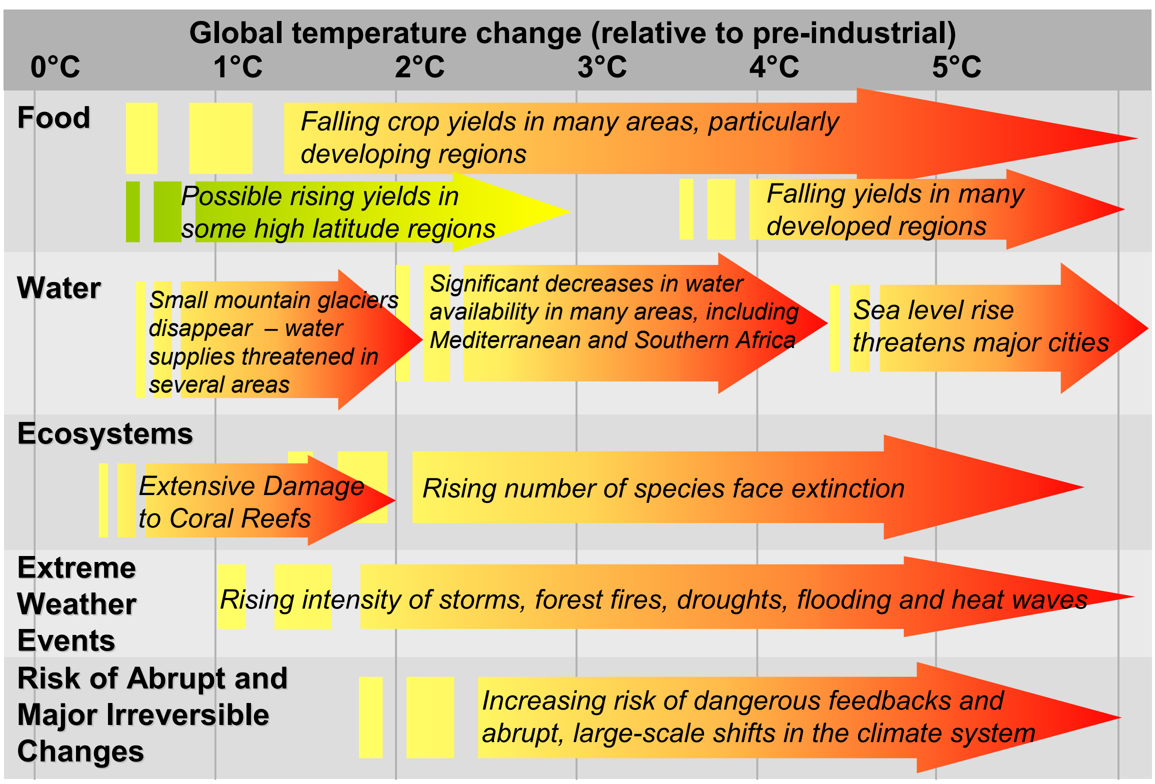

Figure 1: Projected impacts of climate change

It's important to emphasize the need to consider these higher temperatures. Much of the earlier literature—perhaps understandably, given the state of scientific evidence then—focused on temperature increases of only two and three degrees. Those are, of course, serious issues for study, but it is now quite broadly recognized that the analysis must go way beyond that, because temperature rises of four, five degrees and more can happen if the world continues with business as usual.

What about the relationship between the greenhouse gases in the atmosphere and temperature? A key scientific advance of the last four or five years, from the point of view of this type of analysis, is that scientists are now offering us a clearer understanding of the probabilistic relationship between the rise in concentration of greenhouse gases in the atmosphere and the rise in global temperature. Figure 2 illustrates the kind of probability distributions that scientists are now indicating.

Figure 2: Stabilization and eventual change in temperature

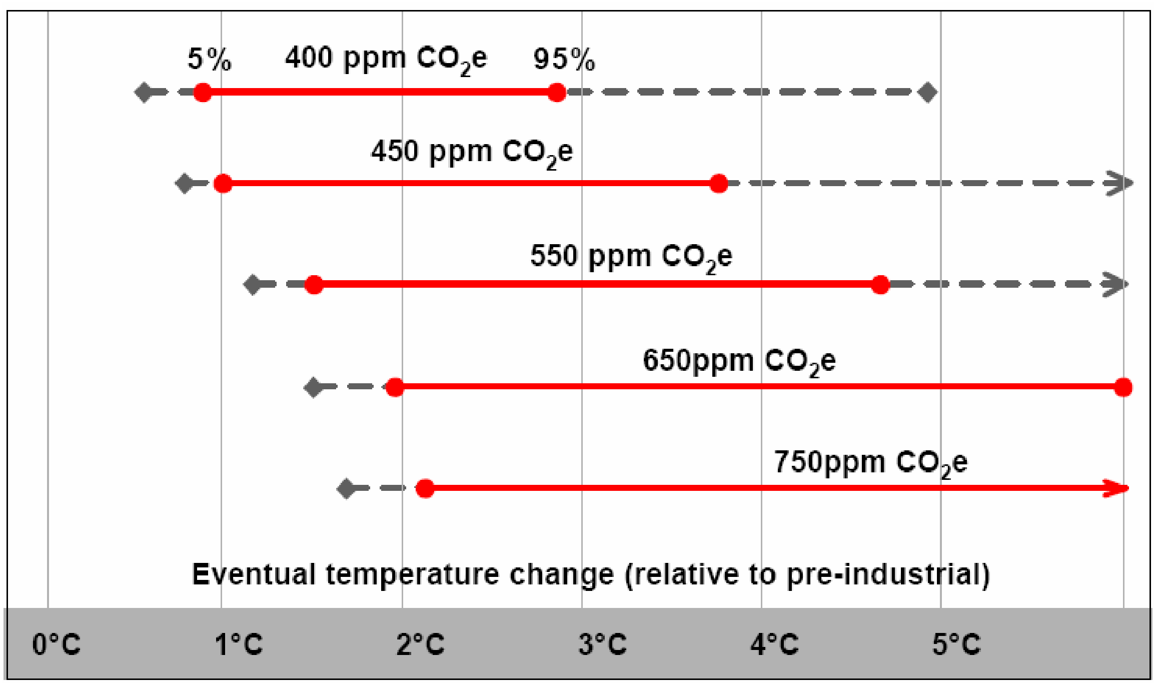

The figure shows eventual temperatures relative to those of pre-industrial times. I emphasize the long lags in the system. The red intervals here are 90 % confidence intervals, so that if you take that particular scientific probability distribution, you have a 5% chance of being off the top end and a 5 % chance of being off the bottom end of that distribution. We drew dotted grey lines there because we have chosen among the possible scientific distributions, and have been moderately conservative in our choice, sticking fairly closely to the kinds of probability distributions used in the fourth report of the IPCC. The dashed lines represent the results from the full range of studies. We drew on the same science that IPCC was drawing on, not as scientists but as consumers of science. And I emphasize this because some authors have suggested we've exaggerated the science in some way. In fact we've just taken science from the same sources as the IPCC and have been fairly conservative as a comparison of the Stern Review and the Fourth Report demonstrate.

Figure 2 thus shows the kinds of risks that the world will be running. We are currently at 430 parts per million (ppm) of carbon dioxide equivalent (CO2e). We're adding about two-and-a-half ppm a year. It's very hard to get carbon dioxide back out of the atmosphere. So, even if that flow of two-and-a-half parts per million a year doesn't speed up, in eight years, we shall be at 450 ppm. But if we do not do much about climate change, that two-and-a-half ppm will go on rising. Though it depends how you model these things, the average rate could easily be four parts per million a year over this century. That would take the greenhouse gas level higher than 800 ppm CO2e, dropping off the scale in Figure 2.

At a greenhouse gas level of 850 ppm CO2e, we have more than a 50 % chance of global warming of more than five degrees Centigrade. So, to talk about the risks associated with temperature increases of five degrees, six degrees, seven degrees Centigrade is not at all fanciful. These are the kinds of risks the world runs if, as a world, we do nothing about curbing emissions.

So, here we need to make a judgment about what policymakers should do. This judgment can be made by comparing the risks we face from different levels of concentration, and paths to stabilization of concentrations, and the costs associated with the paths that stabilize at these levels. Thus we ask, for example, is it worth paying the costs of stabilizing emissions at 550 ppm CO2e to avoid the additional risks of temperature rises beyond this?

What we offer is the suggestion that 550 parts per million of CO_2e is an upper limit on the kinds of risks that it would be sensible to run. Let us be clear, this is a dangerous place to be: 550 ppm gives a 50/50 chance of global warming above or below three degrees Centigrade, and quite a significant probability of its being above four degrees Centigrade. In some models, it gives a small probability of being above five degrees Centigrade. Given how close we are to 450ppm, this represents the most ambitious target that is likely to be feasible. Those who believe that an upper limit of 550ppm C02e is too radical a curb on emissions must come clean and declare that they are ready to accept the very heavy risks involved.

So, this is a view of the stocks and the risks associated with them. After discussing our analytical framework I'll move on to describe the flow paths that would be consistent with stabilization.

Analytical Framework of the Stern Review

Throughout our work on the Review we tried to bring in very strongly the economics of risk, acknowledge uncertainties, flush the ethics out into the open, and see climate change as very much an international action problem.

Modeling is a valuable supplement to the kind of risk analysis that I was just pointing to. But climate change is a very complex and big problem, and when you think of it on the appropriate scale and detail you realize the very simplistic nature of the structure and parameters you can feed in to an overall model. Inevitably the modeling route leaves out a great deal of what's interesting, and in big aggregated world models, whether of climate change or other concerns, you lose the ability to analyze the kind of risks that are run at the regional and at the country level. Of course, when you take the very broad-brush modeling approach, rather than adding up in any direct way these local risks, a tremendous amount starts to turn on the design of your model. So, you always have to be aware of the model structure you are using and the extent to which the kinds of results that come out are influenced by the kinds of things that you put in. In that respect, climate change is a modeling area that's no different from any other. But in this case the magnitude and dimensions of the effects and very long time periods make this problem much more severe than we usually encounter in economics.

Our model choice was strongly influenced by the importance of risk and uncertainty. We chose Chris Hope's PAGE model because for us it was the easiest available model with which to incorporate uncertainty in an explicit way. Second, Chris

Hope was generous in helping us use it. Third, it is valuable in that it is designed to reflect the results that come from other people's models thus ensuring that it is not, in any sense, an outlier.

To assess damages, you have to have a criterion, and the criterion we used—expected utility—is quite standard in the literature on the economics of risk. Damage from climate change is something that takes place over time, so it can be treated as the expectation of a social utility integral. In economics that is a fairly standard way to approach this kind of problem.

To compare damages with costs, you have to calibrate the criterion function in some way. To talk about the sum of social utility or the expectation of the sum of social utility over time in terms of expected utils is not something that most people find it easy to get their heads around. So we used a concept that James Mirrlees and I suggested some 35 years ago: the balanced-growth equivalent (Mirrlees and Stern 1972). In Mirrlees-Stern you calibrate a social utility integral in terms of the initial consumption level that would produce that utility integral if this consumption grew at some standard rate associated with growth in the model. Thus you calibrate the expected social utility integral using the balanced-growth equivalent (BGE) by looking at the initial consumption level, which you think of as growing steadily, and measure gains and losses in terms of that initial consumption level. But, of course, we should think of the gain and loss applying to the whole balanced-growth equivalent path.

This is also a useful tool for incorporating risk analysis as it acts also as a certainty equivalent. We simply calibrate the expected utility integral in terms of the certain initial consumption level growing at the standard rate which produces the expected utility integral. Differences in balanced growth equivalents with and without climate change in this context can be interpreted as the simple annual insurance premiums we'd be prepared to pay to avoid the uncertain losses arising from climate change. They represent the annual premium that we'd be prepared to pay each year to avoid the equivalent paths as a result of climate change.

You can think of the overall expected social utility criterion for damages together with the BGE calibration as embodying three sorts of averaging: averaging over time, through the utility integral including whatever discount factor you attach to it; averaging over space and individuals, because these are effects occurring in many parts of the world and to different people in different ways; and averaging over outcomes. So, when we express the damages that occur from business as usual, relative to a world where climate change does not take place, our estimates involve those three kinds of averages put together. It's important in thinking about them to keep in mind these averaging processes.

How you do the averaging will have a strong effect on the results, and that's one feature of the modeling that you need to think about in interpreting our results in the Review. When we do the averaging, we're discounting, we're treating risk aversion, and we're treating equity. Both the model structure and our judgments about ethics will drive the results. Later I'll discuss the sensitivity of the results in relation to those two things.

Figure 3 shows the equivalent loss in consumption each year from the effects of business as usual climate change. The 90% confidence interval is shown underneath. In this Figure there is a little bit of sensitivity analysis: you bring more things into your account of damages as you move down from the first row to the second column, from very narrow GDP-like estimates (in economic sectors such as agriculture and tourism) to estimates that involve much broader views of impacts (non-market impacts such as on health, deaths and ecosystems). So incorporating a broader set of impacts increases the central estimate of damages from 5% of consumption each year to 11%.

Figure 3: Aggregate estimates of impact

	Market impacts	Broad impacts
Baseline climate	5% Range 0-12%	11% Range 2-27%
High climate	7% Range 1-17%	14% Range 3-32%

Note: The figure shows the magnitude of effects in the middle of the plausible range, taking into account sensitivity analysis in the Review.

The other aspect of the sensitivity here is as you move across the rows, looking at different kinds of probability distributions for the impacts of climate change, moving from 'baseline climate' to 'high climate' introducing in a small way some aspects of the greater uncertainty that science is aware of but has not yet been able to model in a very explicit way. The kinds of probability distributions that the IPCC has calibrated and worked with include only certain things that could happen. These distributions don't say very much, for example, about carbon feedback mechanisms, because IPCC has not yet got a clear understanding of the levels of risk associated with these mechanisms. Scientists don't know enough about them yet, though they know those kinds of effects are there. So in the model we have deliberately left out large areas of uncertainty that have yet to be researched. We did include just one of these possible effects: the melting of the permafrost and the methane release that would result from that. That's reflected in the difference between the first and the second row in Figure 3. It is equivalent to an additional 0.4 degrees of temperature change. The latest IPCC report suggests that these feedbacks could add up to a further degree temperature change by the end of the century. Not surprisingly, you get higher estimates of damages if you include more things in your damage list and if you include broader probability distributions. This is a very important sense in which we have underestimated the damages.

The set of issues around intra-generational income distribution is not represented here, in the formal aggregate modeling because of the time constraints we faced in producing the Review. They are very important to our overall approach however, and are very prominent in our detailed disaggregated analysis. We did speculate informally about the effects of bringing intra-generational distribution into the modeling story. You can do this by relating weighted growth rates to the utility functions that underlie them, and we would suggest that you probably should scale these up by a factor of a quarter or a third. François Bourguignon has done some back-of-the-envelope calculations on that kind of thing, as well, and he gets similar kinds of numbers. So, that's how we got to the 5% to 20% loss in consumption that is often quoted.

Averaging over time is important. Those of you who are not economists, please excuse me while I discuss briefly some of the formal aspects of ethical judgments. We talk about the elasticity of the social marginal utility of income, which we represent by η, and we talk about the pure time discount rate (δ), which you should view as the degree of discrimination between people by date of birth. So, in other words, pure time discounting says that an individual who is born 30 years after another individual would, if you use a pure time discount rate of two per cent per annum, be given half the weight of the individual born earlier. Many people for a century or more have thought hard about the pros and cons of discounting the utility of future generations, and I still have not heard a convincing ethical argument for indulging in that kind of discrimination, particularly in the context of issues which affect the entire planet. We may know lots of people who don't care about the future, but that doesn't mean this is the right ethical standard to apply for such an important issue, profoundly affecting the welfare of future generations.

For the social marginal utility of income, we used as a base case η equals one, which is quite a standard value in the literature. This value means that if somebody is five times better off than somebody else, then an increment in consumption to the person who is better off has a 20% (one over five) weight relative to the person who is worse off. Thus we value an increment in consumption to the poorest at five times an increment in consumption to the richest. Using η equals two would be to say that the person who is five times better off than the other person has, for marginal changes, a one-twenty-fifth (or one over five squared) weight. So, if η equals two, you would be arguing that you could take a dollar from somebody who is earning $150,000 a year and give it to somebody earning only $30,000 a year then even if 95 cents of that dollar disappeared on the way, it still would be a good idea.

A problem in this context is that η is doing more than one job; η is not simply representing egalitarian values but is also a measure of relative risk aversion. This double role might be used to help a judgment as to what plausible social judgments on η might be. However we know from the risk/uncertainty literature that the expected utility model as a description of behavior is problematic. And, in any case, a link from behavior to social values would be a further and difficult step. That doesn't mean that the evidence is irrelevant, but it's difficult to interpret. In a paper we put up on the Web at the beginning of February, we tried to explore some of the philosophical issues and what might be relevant evidence. In this discussion that we must have about ethics, evidence on attitudes to redistribution, attitudes to risk, and attitudes to saving is all relevant and important.

One of the issues that arises with $\eta = 1$ and low δ is that these choices place a big weight on the future. So, it is important to ask "What if you put slightly less weight or a lot less weight on the future? How does it affect the results?" It is important to place this both in the context of sensitivity analysis of the model as a whole and in an understanding of the role of this type of modeling in the overall argument.

Assumptions, Uncertainties, and Results

Why do we get results that are bigger, in many cases, than previous examples in the literature? Is it because we exaggerate the relationship between temperature and damages? I don't think so. The red line in Figure 4 represents the assumptions about impact on global GDP from changes in global mean temperature that were used in Chris Hope's PAGE modeling, and it is roughly in the middle of the range for the literature. Some authors have been more optimistic about the damages likely to be caused by rising temperatures than we were, others less so.

My own view is that the set of assumptions we used understates the damages that rising temperatures will cause to world GDP. It doesn't take into account that at five or six degrees of warming, we're into territory that we really don't understand (our probability distributions account only for 'known' uncertainty not Knightian) and could be very worrying for example in terms of the effects of very large movements of population, including conflict (where it is hard to assess likely aggregate costs). In a moment I'll come to what difference it makes to damage estimates when you assume that damages are more sensitive to temperature.

Figure 4: The PAGE model and other Integrated Assessment models

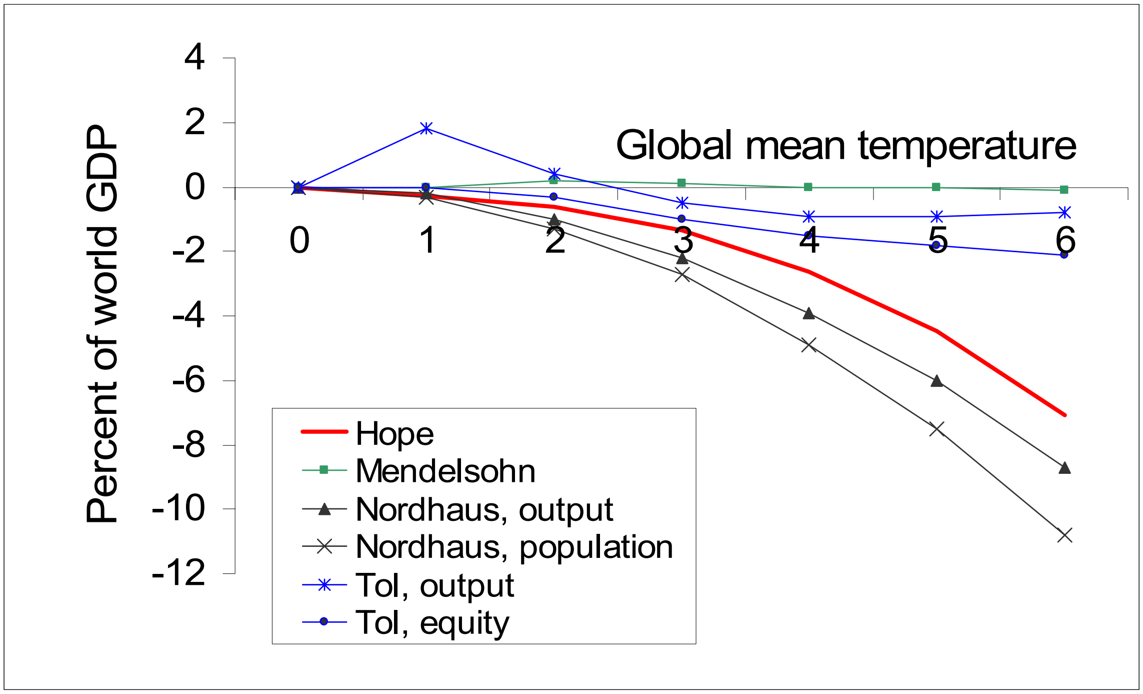

What do we assume about climate sensitivity, the relation between stocks of GHGs in the atmosphere and temperature/climate? The concept of climate sensitivity, which seems to be standard in the literature, measures how much the temperature eventually goes up if you double the amount of greenhouse gases in the atmosphere. Using the PAGE model we look at what happens to temperature if we double the amount of CO_2e in the atmosphere and move from the 280 ppm of pre-industrial times to about 560 ppm. From Figure 5 you can see that the distributions we used are quite cautious but not far off those in the standard literature. The dotted line shown for IPCC refers to the carbon feedback mechanisms I mentioned earlier for which there is some suggestive evidence, but not yet enough to be built fully into the probability distributions.

Figure 5: Estimates of climate sensitivity from Integrated Assessment models compared to Global Climate Models shown in the results of the IPCC and Meinhausen range from eleven studies.

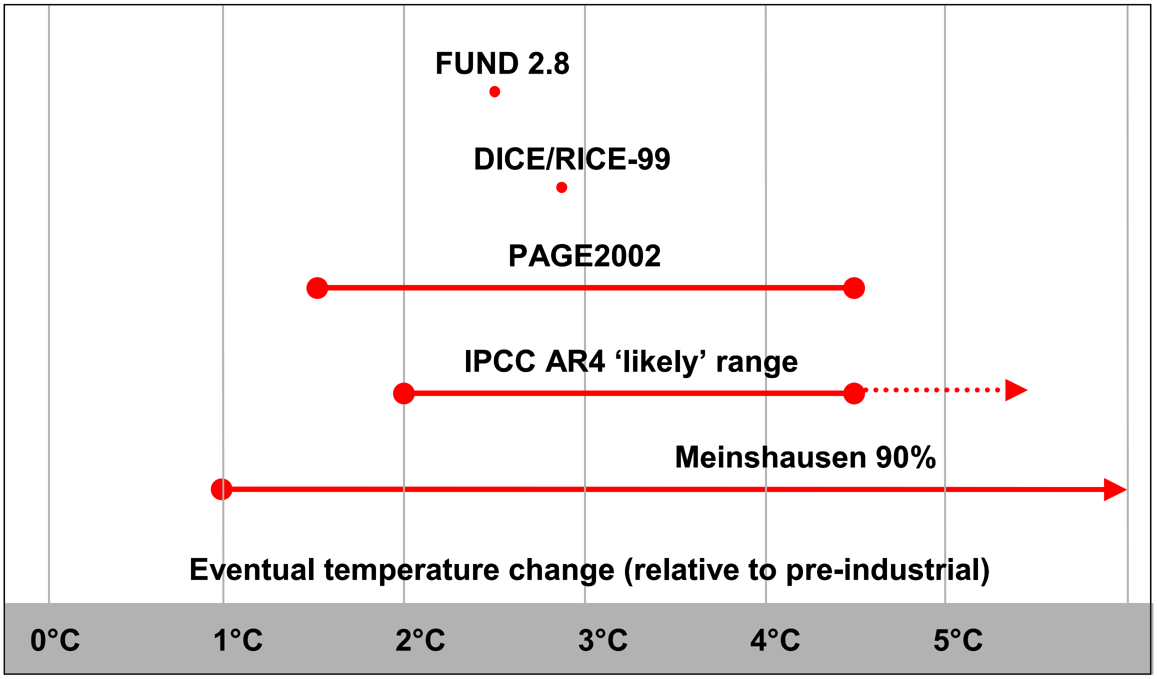

Our assumptions on the relationship between stocks of greenhouse gases and temperature, too, have been within the normal range in the literature. However, we have been very deliberately stochastic here, where some previous studies have not, in my view a crucial omission from these models given the importance of risk to the whole set of issues. These probabilities have become available from the IPCC only in the last few years, but it is crucial now to make use of them in a way that allows us to speak explicitly about the role of risk.

Sensitivity Analysis – What Drives the Results?

Turning now to the sensitivity analysis. Model outputs are driven by model structures and assumptions. We have carried out formal sensitivity analysis, using PAGE2002, the details of which can be found on our website. Figure 6 summarizes the sensitivity of the Review estimates to the four key issues: ethics and discounting; the treatment of risk and uncertainty; adaptation; modeling high-damage scenarios. In each case, the base case in Figure 6, from which deviations are reported, is our 'central' modeling case[1]. The total cost of climate change is derived from a comparison of the 'balanced growth equivalent' or BGE of consumption without climate change to the BGE of consumption after climate damage and adaptation costs have been deducted (see Box 6.3 of the Review)[2].

Figure 6 Sensitivity of total cost of BAU climate change, in terms of a loss in present global mean per-capita consumption (on a BGE path), to various issues.

Variation	Central case	Sensitivity	Change in mean total cost of BAU climate change (percentage points)
Ethical aspects			
Increase in pure rate of time preference, δ	0.1% per year	1.5% per year	-7.8
Increase in elasticity of marginal utility of consumption, η	1	2	-7.5
Model structure			
Failure to incorporate risk and uncertainty	Expected-utility analysis	'Best guess' model based on mode values	-7.6
Increase in relative adaptive capacity of Africa, India and Southeast Asia, and Latin America	Higher and constant relative vulnerability in these regions	Vulnerability instantly falls to that of EU in 2100	-1.5
Increase in damage function exponent, γ	Triangular probability distribution (min=1; mode=1.3; max=3)	3	+23.3
Incorporating recent science	Baseline-climate scenario	High-climate scenario	+3.6
Incorporating risk of 'catastrophic' climatic changes	With risk of catastrophe	Without risk of catastrophe	-2.9

Much attention has focused on our assumptions on value judgments, so it is important to examine the sensitivity of the model results to these. I would recognize that a higher η is tenable especially if you account for responsibility as suggested by Dasgupta. If you increase the elasticity of the marginal utility of consumption, η, as shown in Figure 6, you get two effects. First, because of greater egalitarianism from higher η, if the current generation is less well off than future generations you will put a stronger weight on the welfare of the current generation and less on the damage in the

[1] This comprises the baseline-climate scenario, with market impacts, non-market impacts, and the risk of abrupt, large-scale and discontinuous or 'catastrophic' climatic changes. The pure rate of time preference, δ, is 0.1% p.a., the elasticity of social marginal utility of consumption, η, is 1, the damage function exponent, γ, is sampled from the range 1-3 (mode=1.3), and expected-utility analysis is carried out.

[2] It summarises simulated losses over time, regions of the world and possible states of the world in terms of a permanent loss of global mean per-capita consumption today. In the central modelling case, this loss is around 11% (see Table 6.1 in the Review).

future from climate change. This can make quite a big difference as you move into the future. Second, if you increase the value of η, you also increase aversion to risk, which would affect damage estimates in the opposite way as more damaging outcomes are discounted less heavily and attain a higher expected utility weight. But with the kinds of structural assumptions that we're using, the inequality effect and therefore the stronger discounting effect, dominates the uncertainty effect. With bigger risks, however, the balance can tilt the other way. And these bigger risks may actually become viewed as more appropriate as the science gathers more evidence and analysis.

You can also increase the pure rate of time preference. But, as we emphasized in the Review, if you do that, there will always be a rate of time preference high enough to render future climate effects to have negligible value. If you use pure time discounting of 2% or 3% for a hundred years, you are putting a tiny weight on what happens after a hundred years. If you have a high enough pure time discount rate, neither the economics nor the science matters. As I said, one cannot avoid this ethical discussion, because how one values the welfare of future generations has a direct and very powerful influence on the results. And we were very explicit about that in the Review.

As I noted earlier, we did not formally incorporate a concern for intra-generational income distribution, though we did try to do so in our disaggregated regional work, and with considerable detail. In the modeling work we explored fairly informally and suggested that this pushes up the estimates of damages substantially. Further, a higher η would place still stronger emphasis on the intra-generational distributional issues.

Ethics are important but model estimates, of damages or otherwise, are driven by a lot more than assumptions about ethics. For example, the choice of a lower emissions scenario in the baseline would lower the assessment of the damages, and diminish the case for urgent action. On the other hand, a higher exponent on the damage function, a more comprehensive treatment of uncertainty or allowance for a wider world income distribution (for most of the modeling there is simply aggregate consumption and consumption per head) – all of which are plausible - would raise the projected impacts potentially quite significantly. In fact, it is notable that the degree of convexity of the damage function is perhaps more important than anything else is driving the aggregated results.

My own view based on discussions with leading scientists, particularly thinking about warming of five, six degrees, or seven degrees Centigrade, is that we may have underestimated the degree of convexity, the strength of the relationship between temperature and damages, or at least we ought to allow for it being stronger. In other words for the higher temperatures damages may rise far more quickly with rises in temperature than in the base case used in the model here. You can investigate this stochastically: you can say "we're not sure what sort of convexity to put in", and you can make that parameter one of the Monte Carlo parameters that Chris Hope uses, (and it is indeed a Monte Carlo parameter). We found, not surprisingly, that when we assume a stronger relationship between temperature and damages, the model results show much larger damages.

In addition to the issues we have discussed in detail, there are many other assumptions and judgments that will affect results. We summarize a selection of these in Figure 7, along with a back-of-the-envelope calculation of how they will affect our estimates. For example, if economic growth is faster than we assumed, you get sharper discounting, because future generations will be richer, but you also get emissions going up more quickly. So, those two things act in different ways, but our belief is that, net, faster growth would probably increase the damages estimated here (line 3, Figure 7).

Figure 7. Further sensitivity of total cost of BAU climate change to various assumptions.

Variation	Central case	Sensitivity	Change in mean total cost of BAU climate change (percentage points)
Ethical aspects			
Accounting for intragenerational income distribution/equity weighting	Not included	Included	+6
Population growth	IPCC SRES A2 scenario, extrapolated by Hope (2006), gives global population of 21.5bn in 2200	Reduce population growth by 40% over modeling horizon, whilst holding emissions constant.	-4
Model structure			
Output growth	200-year average of 1.3% per capita	Increase annual per-capita growth by 1%	+
Terminal conditions	Modeling horizon ends in 2200, emissions fall instantly to a rate equal to the Earth's natural capacity to absorb GHGs, allowing the impacts of climate change to stabilize	Continued emissions growth post-2200	High sensitivity ++
Aversion to irreversibilities and ambiguity	None	Included	+
Rise in the relative price of environmental goods compared with other consumption goods	Utility is only an aggregate function of total consumption	Utility is a function of both consumption and environmental goods and services	+2
Inclusion of 'socially contingent' risks, e.g. conflict and migration	Not included	Included	++

Commentators have rightly pointed out that our results are sensitive to terminal conditions. We've assumed that damages fall away after 2200, or at least stay constant (in percentage terms) in a world where business as usual emissions continue to rise unabated. Of course, if we allowed damages to continue growing beyond that date, our estimates would increase (line 4, Figure 7). In such circumstances, further areas of risk and uncertainty are left out in addition to the omitted carbon feedbacks already noted. Nor have we accounted for aversion to irreversibilities and ambiguity (unknown possibilities), that is, risks that we are unable quantify, for example those associated with pushing temperature into uncharted territory.

We emphasized that there may be more damaging feedbacks and impacts at high temperatures that we have not yet identified. We explained in Chapter 2 of our review (and referred to the work of Claude Henry) that the Knightian distinction between risk (known probabilities) and uncertainty (unknown probabilities) is likely to be important here. Whilst exploring this issue and pointing to relevant analytic techniques, we have not tried to estimate the consumption we might be willing to forgo to avoid this ambiguity.

So, I think there is actually more weight in the tail of the damage distribution than is represented in our results. Except for that concerning population size, most of our other sensitivity analyses would point to higher estimates than ours. So, on the whole, on the structural side of these models, we've been probably very conservative. That is to say, from the sensitivity analyses a wide range of possible estimates emerges, but in most cases the relevant variations in structural assumptions would increase our damage estimates.

Figure 8. Explanation of further sensitivities

Accounting for intragenerational income distribution/equity weighting	In the Review, we did not have the opportunity to model the regional impacts of climate change. Given a positive elasticity of marginal utility of consumption, consistent valuation of the impacts of climate change across time, risks and regions of the world implies that consumption effects in poorer regions of the world should receive higher weight, just as increments in global consumption today should be weighted higher than increments in global consumption in the future, if the future is richer.
Population growth	Where population growth is exogenous, the social welfare function is weighted by the total size of the population. In Chapter 6 of the Review, we used an extrapolated version of the A2 scenario from the IPCC's SRES (Nakicenovic and Swart, 2000; extrapolated by Hope, 2006) to project GHG emissions, output and population growth.[3] Although the A2 scenario appears, on current trends, to predict a sensible path for GHG emissions, it forecasts a very high global population, reaching around 21.5 billion people in 2200 (as extrapolated by Hope, 2006). As a result, the cost of climate change will be higher than it otherwise would have been, all else equal, because high per-capita costs of climate change next century are multiplied by a high global population.
Output growth	A change in output growth will produce an ambiguous result. Higher annual growth will result in higher emissions. Given the close relationship between output and emissions, a 1% increase in annual growth would likely raise the atmospheric stock of GHGs by a factor of 3 or 4 by early next century, in turn probably quadrupling climate impacts by then. On the other hand, the average annual consumption discount rate would increase by 1 percentage point, before climate impacts. The effect is likely to be finely balanced at first, but reasonable assumptions suggest that steeply rising climate damages, brought about by such a high stock of atmospheric GHGs, dominate over the longer term[4].
Terminal conditions	In other words the length of the modeling horizon and what is assumed to occur thereafter. The PAGE2002 modeling horizon runs until 2200. Thereafter, the Review in effect assumes that emissions fall instantly to a rate equal to the Earth's natural capacity to absorb GHGs, allowing the impacts of climate change to stabilize and the stock of GHGs to rise very slowly. The longer the modeling horizon, the higher the costs of climate change, though in the very longest run, the coupled climate-economy system may eventually regulate itself, even in the absence of policy.
Aversion to irreversibilities and ambiguity	We did not explicitly account for aversion to having to make irreversible decisions – the number of such decisions is likely to increase in line with the stock of GHGs, adding further to costs. In addition, we did not formally take account of ambiguity aversion, which becomes important

[3] Our regional growth rates were converted from market exchange rates to purchasing power parities.

[4] Output growth would also affect adaptive capacity and willingness to pay to avoid non-market climate impacts. Again the effect is ambiguous. We deal with adaptation separately, but note that willingness to pay to avoid non-market impacts is a quantitatively important component of most IAMs that include them (see Warren *et al*., 2006).

	where the consequences of climate change cannot adequately be represented by a continuous probability distribution.
Rise in the relative price of environmental goods compared with other consumption goods	We can expect the relative price of environmental goods to rise compared with other consumption goods, but this is not captured in a utility function where aggregate consumption is the only numéraire. Thus climate impacts are likely to be underestimated (e.g. Tol, 1994).
Inclusion of 'socially contingent' risks, e.g. conflict and migration	No IAMs yet take explicit account of socially contingent costs, which would increase damage estimates.

The damage-impact estimates in the Stern Review are higher than some of the existing literature for three reasons. First we used the latest science in 2005/06 and this suggests larger temperature changes than the previous studies (which reflected the latest science at their time, many have data/estimates from the late nineties). Secondly we have explicitly included in the analytics of the economics of risk the latest probabilistic assessments: climate science now is much more specific on probabilities. Finally, as discussed at length today we have dealt with the ethics explicitly and have argued that the values likely to emerge from this more explicit discussion would increase the weight given to the future and thus estimates of damage compared with some of the more casual ethical treatment of the earlier literature. While most attention has focused on the last of these revisions, all three are very important to our results. Indeed one can argue that for plausible parameter values each had a similar impact on the estimates.

We should also explain key elements from Chapter 13, the summary chapter for the first half of the report which puts together the preceding chapters including the modeling chapter (Chapter 6). Many commentators seemed to have overlooked this in their decision to focus on Chapter 6. It is very important to see the sequencing of the logic here.

In Chapter 13, inter alia, we explain that the key question of the first half of the Review is whether we would pay 1% of GDP to avoid the risks of damages discussed in Chapters 1, 3, 4 and 5 (with multidimensional regional flows). Thus chapter 6 is a useful but supplementary analysis. And we stress that stabilization between 450 and 550ppm CO_2e would avoid most of the damages from climate change (around 90%) but not all.

Overall, from this discussion of sensitivity, there seems little justification in changing our broad view that the cost of avoidable climate change, (i.e. comparing a 'curbed' emissions path with BAU) is well in excess of the cost of stabilizing GHG emissions in the range of 450-550 ppm CO_2e. We did not undertake an explicit optimization, data constraints and uncertainties make this too ambitious. Those that have taken this route rely on ad-hoc assumptions, to which results are likely to be very sensitive. But in Chapter 13 we do proxy an optimization process to determine the likely stabilization range. Hence we summarized the literature which shows that the incremental raising of emissions above 550 ppm CO_2e adds more to damages than it costs to reduce emissions by a ton of carbon.

It is very clear that IAMs produce results that are very sensitive to assumptions, which is why they should be used with caution. Indeed, it is clear from the sensitivity analysis that in many respects we were cautious about, or omitted, many aspects of the modeling structure, which would have raised damages. There is certainly no justification in the claim that we systematically chose assumptions, which would give high damages.

It is very important to see the sequencing of the logic here. We start from the economics of risk, by emphasizing the importance of a quantitative stabilization target, in this case we identified a range. If this is no higher than 550ppm CO_2e, then there are strong implications for the emissions time path. Whilst there are some timing choices such a path should peak within 20 years and be below 30% relative to current levels (in absolute terms of 2050). We must be very clear on this. The role of the price mechanisms then become the decentralization of decisions to keep down costs around such an emissions path. In this sense the argument goes from quantities to decentralization prices. It does not, and should not start with the social cost of carbon (SCC). The SCC is a useful concept but is very sensitive to assumptions, indeed far more so than the stabilization target.

Research Questions on Costs and Damages

Many people will create their own models, building in risk and uncertainty in different ways, and with different approaches to ethics. What I think is particularly important is to put a lot more weight into understanding the probabilities and impacts of temperature increases and the tails of the distributions, either through using a more convex function for damages or through expanding the distributions—and we've got powerful evidence for thinking we have understated the spread. The relationship between temperature and damages is something that is a bit more speculative, but given the kinds of disruption that could occur at five, six, seven degrees Centigrade, I think that using a more convex damage function has considerable plausibility.

Other research questions needing further work include:

- Exploring ethics, including disentangling inter-temporal distribution, intra-generational distribution, and risk aversion. I think that using one parameter to represent all these things is much too crude.

- More detailed regional and sectoral analyses of both adaptation and mitigation. It is this detailed analysis of risks and action that should provide the basic underpinning for the analysis of action on climate change.

Policy Issues

Figure 9 emphasizes the urgency of action on climate change. The yellow line shows that if you adopt the target of stabilizing CO_2e concentrations at or below 550 ppm, then you should start cutting back, in absolute terms, the world's emissions in the next 20 years and keep them falling fairly steadily thereafter. To stabilize below 450ppm CO_2e would require emissions to peak by 2010, with a 6 – 10 %annual decline thereafter.

Figure 9: Emissions paths to stabilization

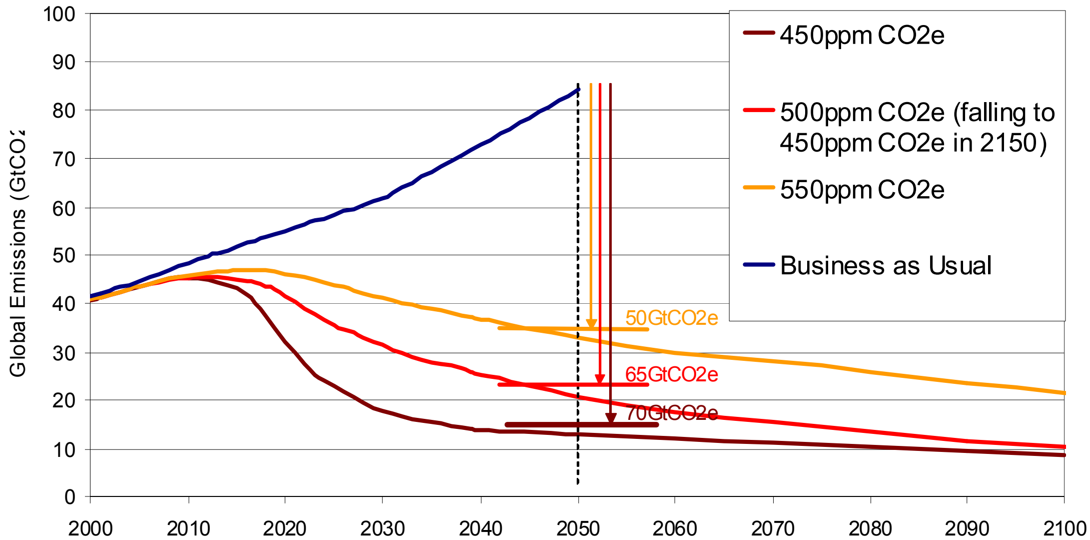

Many people are angry with us for even suggesting that 550 ppm is a sensible target; they would argue very strongly that 500 or even 450 is the maximum one should accept. They have a case and I gave the arguments on that before. All these targets clearly involve strong action.

Those that claim our implied action plans are too radical should be quite explicit that they would propose a path which goes considerably above 550 ppm CO_2e and would accept the corresponding risks involved. Unfortunately many have failed to be transparent on this by burying their arguments in the level of the carbon price.

What would be the costs of action to stabilize below 550 ppm CO_2e? We estimate the cost of a trajectory consistent with stabilization at 550 CO_2e to be 1% of GDP. We made our calculations in two ways: by surveying the results from the various macro models (suggesting a range of plus and minus 3%of GDP in 2050 around the central estimate of 1%), and by looking at the resource costs of particular kinds of technologies (suggesting a range of minus one to plus 3.5%of GDP in 2050). Shortly after we published these results, the International Energy Agency published a much more detailed study than we could possibly manage, whose cost estimate was slightly lower than ours (IEA, 2006). I think that one %is a reasonable ballpark for a 550 ppm CO_2e stabilization

Given the costs of the impacts of climate change, taking urgent action is good economics. If you compare the 1% estimate for the costs with the kinds of damage that

I've been discussing, then taking strong action to stabilize at or below 550 ppm looks like a very good deal. It is hard to argue that a 1% increase, on average, in cost is going to make a big difference to the competitiveness of the United States or Europe versus China or India. Those are the kinds of cost differences that people absorb and just get on with during a growth process. And, of course, there will be opportunities, and there could even be Schumpeterian kinds of growth surges on the back of new technologies. There are potential important co-benefits associated with energy security, cleaner air and so on, as well.

The issues around estimating the costs of mitigation are set out in Chapter 9 and 10 of the Review. Time constraints do not permit me to elaborate these further in this lecture but I would like to emphasize our estimates are based on the existing mainstream literature and supported by our additional research.

Key Principles of Policy

Two-thirds of this lecture has been on costs and damages. More time would have allowed me to 'correct this imbalance'. The report is roughly 50-50 on costs–damages and policy and much of the more subtle and difficult economics lies in the policy analysis.

In considering policy we should be pricing for externality. That's principle number one, basic, absolutely right, and terribly important. As well as proper pricing for carbon, we should be promoting research, development, and deployment. There's also an important discussion to be had about behavior of individuals, firms and governments beyond simply the appropriate price incentives.

Carbon Pricing

Appropriate price signals for carbon can be established in different ways: greenhouse gas taxes, capping emissions and setting up a market in permits, or implicitly through regulation. In thinking about policy instruments we have to think about the risks. Here the risks of overshooting concentrations in the medium term are very high. Thus, the economics of risk points you to first thinking about stabilization stocks and then about the flows of emissions that are consistent with those stocks. Then as we think about managing adjustment costs, we go on to think about price mechanisms.

As we noted in the Review, taxation, emissions trading and regulation can all deliver a price signal for carbon.

Different countries will choose different combinations of these approaches for different sectors, reflecting their existing policy mix, histories, conditions and national politics. To take the case of taxes, they may be most useful in pricing carbon emissions from sectors that have a large number of small emission sources, which may also be mobile (such as road vehicles). In such sectors, the transaction costs for a large number of small emitters being involved in emissions trading schemes may be prohibitive.

In other sectors that have large, stationary sources of emissions (such as electricity generation or heavy industry), transaction costs for involvement in trading will be lower, making them more suited to using emissions markets. In many cases, these sectors are also competing internationally. Inclusion in an international trading scheme therefore helps to reduce any risks of differing carbon prices being imposed at the domestic level that have may impact on competitiveness. In theory, taxes could be harmonized across countries to prevent any competitiveness impacts, but experience in other policy areas shows this is very difficult to attain.

In terms of impacts on international co-operation, trading has the advantage that it opens up markets for emissions reductions across borders and therefore allows automatically for transfers of finance and investment between countries. Access to broad international markets is likely to allow firms to locate least-cost options for reductions and therefore keep compliance cost low. Where developing countries are involved in such markets, it therefore offers a channel for the financing of low carbon investments in these countries, which is particularly important for international co-operation on climate change. Again, in theory, carbon taxation could be used to transfer revenues across borders, but in practice this would be more difficult to achieve than through the direct, and largely private sector transactions that occur within an international emissions market.

Expanding and linking the growing number of emissions trading schemes around the world is a powerful way to promote cost-effective reductions in emissions and to bring forward action in developing countries: strong targets in rich countries could drive flows amounting to tens of billions of dollars each year to support the transition to low-carbon development paths in poorer parts of the world.

The economics of cost points to a need for short-run flexibility within sectors and countries. Policymakers and markets should be able to respond to new information on impacts and costs. Credibility, flexibility, and predictability are key if policy is to influence investment.

Research, Development, and Deployment

Some would say then that if you fix the price, if you fix the externality, there is nothing else to do; just let the markets work including for R&D, innovation and deployment. But I think that in the case of climate change this is misleading. For example, the markets would never be totally confident about a pricing policy based on the entire world acting together over the next few decades. So I think we have to go beyond that. And we know that in a world full of market imperfections and constraints on taxation, simply pricing for an externality in terms of marginal cost will not generally be the best policy.

One element of policy to promote technological change is public funding to support innovation in new technologies. Global public R&D in energy has been cut by half in the last 25 years or so. We can't say for sure what is the right level of global R&D, but such a large cutback doesn't sound sensible in relation to the climate trend. So, in the Review we suggest that R&D should at least be doubled, back to around $20 billion annually. Incentives for deployment should increase two to five times, from current levels of $34 billion. As Figure 10 shows, cumulative experience brings down costs. It is important to develop a portfolio of mitigation technologies, as this will reduce mitigation costs in the longer run. Governments must ensure that they provide adequate incentives to spur the development of what are likely to be the key technologies of the future such as carbon capture and storage for hydrocarbons. This is of particular importance for coal which is not only the most polluting hydrocarbon in relation to climate change but is also by far the most important source of energy for electricity generation, and is likely to remain so, for many countries including China and India.

Figure 10: Technology needs more than a carbon price (Source IEA, 2000)

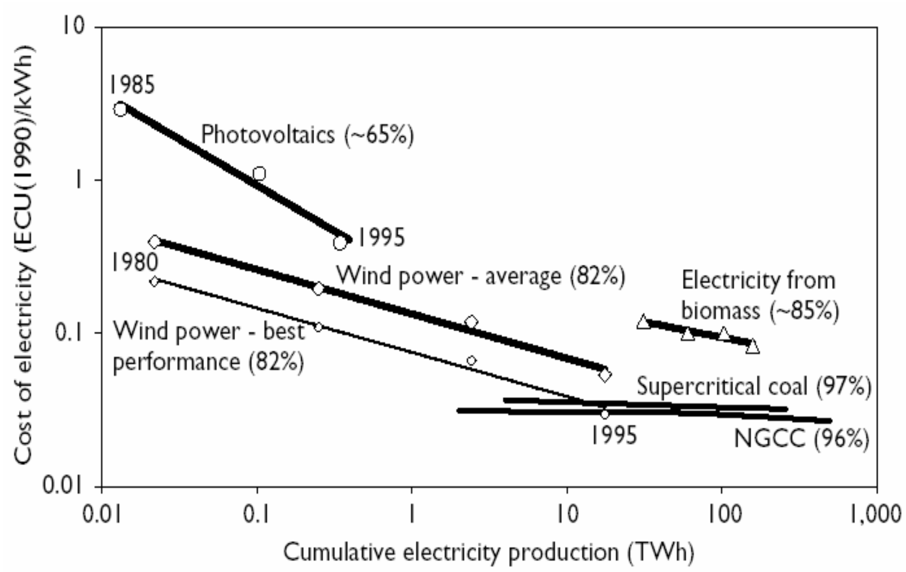

Other Market Failures and Behaviour Change

People will want to discuss what responsible behavior is, just as they discuss what responsible behavior is on recycling, without necessarily being totally influenced in their choices only by the relevant prices and costs. And there's a discussion to be had about related market failures in buildings or landlord/tenant relationships, such as "Will you capture the cost of investment in insulation and so on?" There are other relevant kinds of market failures, particularly as regards energy efficiency.

As I've already emphasized, *responsible climate change policy* can be consistent with *growth* and *energy security*. If, but only if, we design our policies well, those things can be brought together. This is a crucial insight for policy and has begun to be very important in discussions on both sides of the Atlantic. Demonstrating this in greater detail at country and regional level will be crucial in taking policy forward.

International Action

Climate change is an international collective action problem. Such action requires a common understanding of the scale and nature of the problem. It also needs transparency and mutual understanding of actions. International institutions, including through their surveillance, can have a key role on of both these.

Partly this means being able to recognize what other countries are doing. For example, it is important to understand the advances in energy efficiency the Chinese are making in their 11th Five Year plan. Yet, they're still doing a lot of things that are very polluting, including the opening of 1 or 2 coal-fired electric power stations a week. But they're also doing a lot of things, such as reforestation, that go the other way. In China, we take care to explain what's going on in the U.S., and the U.S. is not inactive in this area, and in the U.S., we're keen to insist that China is doing quite a lot, too. Indeed, there are signs that most countries are starting to get to grips with the problem, including India in its 11th Five Year plan, which is starting this year. However, the scale of action in most countries is still far too low.

We should also be looking for trading structures that sustain cooperation by giving people gains from coming together. This is one of the great advantages of cap and trade.

Trading structures must be equitable, too, because there's a very powerful feeling in the developing countries, which many of us would share, that "rich countries put those GHGs there and now they're telling us it's time to slow down our growth." Well, first we have to argue that growth doesn't need to slow down. You can move to low-carbon growth without necessarily slowing down. There will, however, be some associated costs, and it's only right that the rich countries take on the bulk of these costs given their historical responsibility, their wealth and their access to technology.

So, if the overall target should be at least a 30 %overall cut in emissions by 2050 on a global scale, the range of 60 to 90 %in rich countries seems about right from the equity point of view (see Part VI of the Review). That's California's 80 %, that's France's 75 %, that's the UK's 60 %, and so on. Effective international action needs to go ahead on several fronts: building carbon markets; deforestation; and adaptation, including through development aid, sharing of new technologies, and support for international public goods.

Building Carbon Markets

The demand side in world carbon markets comes from strong ambition in the rich countries. On the supply side in those markets we must try to identify more clearly than we've done up to now what constitutes an emissions reduction. We need better benchmarks, and we need stronger institutional structures to support these markets. Figure 11 shows how carbon markets could be scaled up from where we are now. In the bottom left-hand corner in the light yellow we have the existing European Union emissions trading scheme. The darker block on top shows the potential size if expanded to all sectors. Moving to the right shows how markets could be expanded if other regions were involved.

Figure 11: Possibilities for building international carbon markets

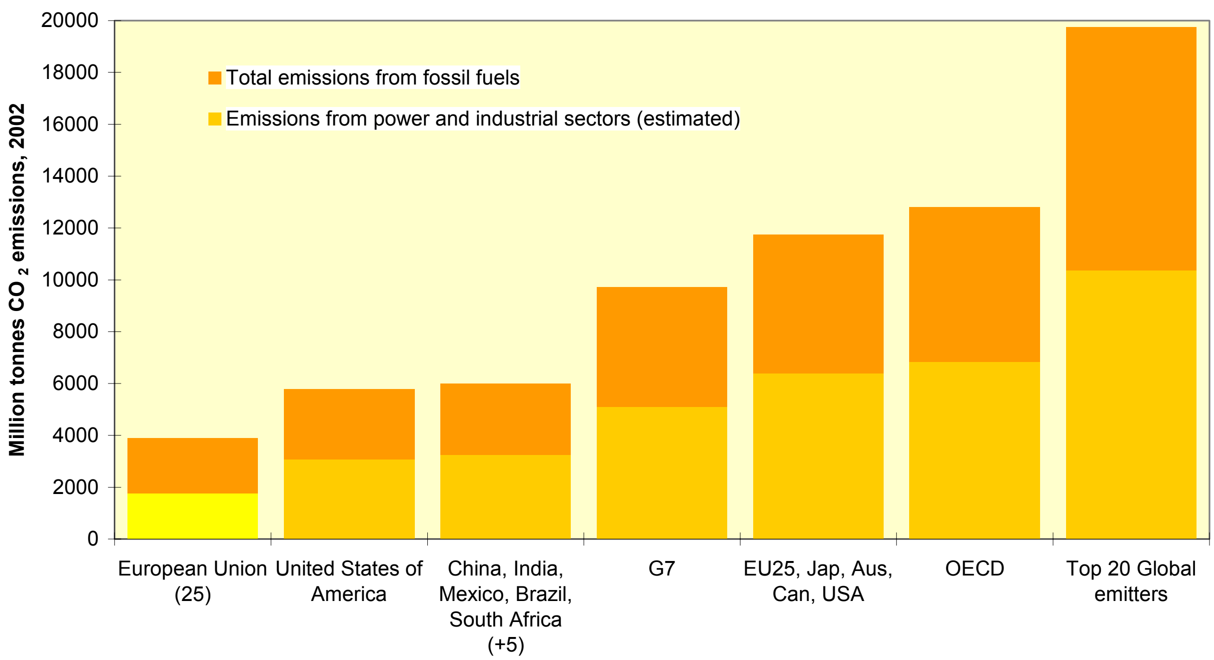

Deforestation

Figure 12 shows the need to look more broadly than just at electric power or transport. You have to look across the board at industry and buildings and, of course, land use where the key issue is deforestation. There's some dispute about the numbers, but deforestation currently looks to be a more important source of greenhouse gas emissions than transport. So moving ahead strongly to curb deforestation could be highly cost-effective and significant. Forest management should be shaped and led by the nation where the forest stands. Large-scale pilot schemes could explore alternative approaches for providing effective international support.

Figure 12: Sources of emissions, by sector

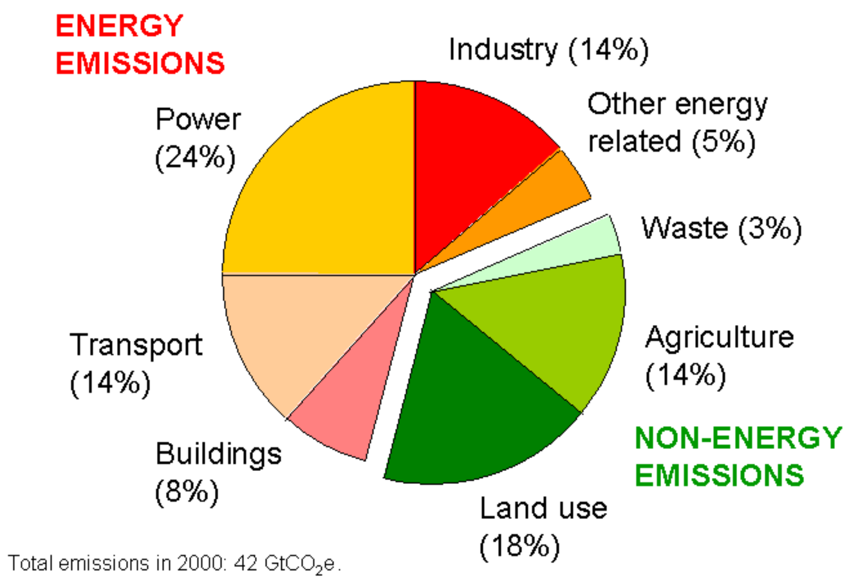

Actions on deforestation, as for any other sector, only take you part of the way. We are also going to have to look at energy efficiency in electric power, transport, buildings, industry, and so on. Reducing emissions wherever cheapest will minimize mitigation costs. So emissions should decline much more in some sectors where cheaper than others. Thus there are benefits from designing policy to ensure that there is a similar price in different sectors and regions. Increased energy efficiency and combating deforestation are the fastest and cheapest ways of reducing emissions in the short run.

Adaptation

Climate change will require costly adaptation in all countries. But adaptation to climate change will put strong pressure on developing countries' budgets and on development aid. Climate change adds to the complexity of decision-making about aid budgets, the productivity of aid, the need for aid, and the budgetary challenges of developing countries. If the environment gets more hostile, then the cost will rise of building bridges and railways and roads and irrigation systems to withstand the greater threats that they will face. We don't really know yet how much the extra cost will be, but it could be tens of billions a year for developing countries, and it surely underlines the arguments for delivering on the aid commitments that were made in Monterey in 2002, the European Union in 2005, and the Gleneagles G8 summit in 2005—especially when you couple it with the equity and the historical responsibility arguments on climate change. I did think of recommending more official development assistance in the Review, but having spent, with many others, 30 years doing that and seeing very strong promises in 2005 in the EU and at the G8, I decided we would just say, "The argument we gave you then was very powerful, and now it's overwhelming, so deliver on your promises."

International action also has a key role to play in supporting global public goods for adaptation. Many international public goods are relevant here, including more climate-resistant crops and technologies, and disaster responses. Weather and climate-change forecasting, of course, are of great importance in any kind of adaptation.

Research Questions for Policy

Research questions relevant for policy include:

- How to link and expand emissions trading schemes. There are lots of technicalities here that matter.

- How to develop and deploy carbon capture and other key technologies globally. For example, I have not emphasized carbon capture and storage of coal much in this talk, but, around the world, approximately a half of the electric power comes from coal. The proportion will be 70 %or more from coal in India and China for the next 20 or 30 years. Unless we get to grips with more efficient coal and carbon capture and storage for coal, we'll be ignoring a big part of the problem. This is the only technology that we explicitly highlight in the Review. Otherwise, we suggest that countries should adopt the right kind of incentives, and see what kinds of technologies emerge.

- Planning for adaptation. Adaptation relies on a great deal of information about the challenges at the local level.

Conclusion

To sum up, the Review concludes that:

- Unless greenhouse gas emissions are curbed, climate change will bring high costs for human development, economies, and the environment. Concentrations of 550ppm CO_2e and above are associated with very high risks of serious economic impacts. Concentrations of 450ppm CO_2e or below will be extremely difficult to achieve, given where we are now and given current and foreseeable technology.

- The costs of mitigation are modest relative to the costs of inaction and strong mitigation is consistent with economic growth.

- Strong international action is urgent: delay means greater risks and higher costs. International action should be designed in a way that is equitable in the international division of responsibility—given the past history and the "common and differentiated responsibilities", in the language of Kyoto—and efficient, working through markets.

- Even with strong action to reduce greenhouse gas emissions, substantial adaptation is essential both for rich and poor countries.

The Review argues for using market mechanisms to address a market failure, taking into account risk in a very direct way when we think how to use those market mechanisms. Thus we should start with a stabilization goal, look at the path to stabilization, and then apply market mechanisms to different sectors and regions where market conditions vary and different incentives are important. R&D should be pushed forward for the usual kinds of arguments, but in this case, there's greater urgency and greater doubt as to whether price mechanisms alone are enough to achieve what's needed.

In the Review we suggest that initially the best path to reducing worldwide emissions is for individual economies in the rich world to take responsibility for strong reductions. If, as they are, California takes on the responsibility, and the EU takes on the responsibility, and they decide to see how they can trade emissions, markets will develop. It will be the insistence of the people of those countries and regions that their governments act responsibly that will be the enforcement mechanism. As individual countries follow their lead and build up their responsibilities, we could start to seek a stronger international treaty, but the world does not necessarily need some grand international enforcement mechanism. Though we certainly don't want to rule this out, it is quite hard to think how such enforcement would really work. But a treaty can be a very valuable signal of long-term international commitment, and I believe, on the basis of the foundations of commitments in individual countries and regions, that it will be a crucial element of international action.

Going back to the argument of why you'd want to stabilize emissions at 550 parts per million, and recognizing we're already at 430 parts per million and adding two-and-a-half parts per million a year and rising, then I think the case for urgent action is very powerful. The later we leave it, the more difficult it will be. The more you ramp up the stock of greenhouse gases in the atmosphere, the bigger the risks you run, and the tougher it will be to get to a sensible stabilization goal.

References

International Energy Agency (2000): 'Experience curves for energy technology policy', Paris: OECD/IEA.

International Energy Agency (2006): 'World Energy Outlook 2006' Paris: OECD.

Intergovernmental Panel on Climate Change (IPCC) (2007). *Fourth Assessment Report of the Intergovernmental Panel on Climate Change.* Cambridge: Cambridge University Press.

Mirrlees, J. A., and N. H. Stern (1972). "Fairly good plans," *Journal of Economic Theory* 4 (2): 268-288