REFLECTIONS ON THE STERN REVIEW (1)

A Robust Case for Strong Action to Reduce the Risks of Climate Change

Simon Dietz, Chris Hope, Nicholas Stern & Dimitri Zenghelis

1. Introduction

Those who deny the importance of strong and urgent action on climate change essentially offer one of, or a combination of, the following arguments. First, there are those who deny the scientific link between human activities and global warming; most people, and the vast majority of scientists, would find that untenable given the weight of evidence. Second, there are those who, while accepting the science of anthropogenic climate change, argue that the human species is very adaptable and can make itself comfortable whatever the climatic consequences; given the scale of the outcomes that we now have to regard as possible or likely under business-as-usual (BAU), this must be regarded as reckless.[1] Finally, there are those who accept the science of climate change and the likelihood that it will inflict heavy costs, but simply do not care much for what happens in

Simon Dietz is Lecturer in Environmental Policy at the London School of Economics and Political Science, and a Visiting Fellow at the Tyndall Centre for Climate Change Research. **Chris Hope** is Director of the Management Studies Tripos and University Senior Lecturer in Operational Research at the Judge Business School, University of Cambridge. **Nicholas Stern** is Head of the Stern Review of the Economics of Climate Change, Adviser to HM Government on the Economics of Climate Change and Development, and Head of the Government Economic Service. From June 2007, he will take up a new appointment as the first holder of the IG Patel Chair in Economics and Government at the London School of Economics. **Dimitri Zenghelis** is an Economic Advisor to the Stern Review of the Economics of Climate Change, HM Treasury.

[1] The scientific evidence on these two points has been summarised, for example, in Schellnhuber *et al.* (2006), and IPCC (2007).

the future beyond the next few decades; most would regard this as unethical. This paper deals primarily with the latter two arguments. Policy issues, especially at the international level, are discussed in a companion paper, in this issue.[2]

A basic conclusion of the Stern Review is that the costs of strong and urgent action on climate change will be less than the costs thereby avoided of the impacts of climate change under BAU. A number of commentaries on the Review have challenged elements of the analysis that lead to this conclusion. This paper sets out our response to these commentaries, including three that were published in the previous issue of *World Economics* (Carter *et al.*, 2006; Byatt *et al.*, 2006; Tol and Yohe, 2006). A second purpose is to further the analytical discussion and to look to future research. We demonstrate that the conclusions of the Review are robust—in particular that the costs of action are much less than the costs of inaction—and do not rest on any one particular modelling approach or assumption. Many of the assertions about our conclusions are simply wrong, and many of the most common mistakes are detailed in the appendix. It is surprising and regrettable that, for example, Byatt *et al.* (2006) and Tol and Yohe (2006) feel able to make such strong assertions on the basis of analysis that is so confused.

We stand accused of over-pessimism in our assessment of the impacts of climate change, and over-optimism in our assessment of the costs of greenhouse gas (GHG) emission reductions. Yet our estimates on both sides of the ledger are drawn from a comprehensive and up-to-date sweep of the literature. Our conclusions on future climate change have, since publication, been affirmed by those of the Intergovernmental Panel on Climate Change (IPCC, 2007), while our assessment of the likely costs of mitigation has received the approval of the International Energy Agency (IEA).

The Review builds its assessment of the costs of inaction—or, put another way, the benefits of climate-change mitigation—first and foremost around a disaggregated analysis of physical impacts worldwide, most prominently on temperatures, the water cycle, sea-level rise, and extreme events. The economics of climate change rests neither solely nor even

[2] Lorraine Hamid, Nicholas Stern, and Chris Taylor (Stern Review team, HM Treasury): 'Reflections on the Stern Review (2): A Growing International Opportunity to Move Strongly on Climate Change', *World Economics*, 8 (1): 169–186.

primarily on the aggregate valuation of climate-change impacts, which comes from (inevitably) grossly simplistic aggregate economic modelling over centuries. Thus many critiques are fundamentally misplaced in their emphasis. The aggregate modelling exercise in the Review was supplementary to the wider disaggregated statement of risks. We argued clearly and strongly that aggregate studies must be treated with great circumspection. Nevertheless, they do play a valuable role in explaining the logic of important issues and that is why we paid them attention.

But to elevate these models to centre stage is to badly misunderstand their potential role in policy analysis and decision-making. It is hard to imagine how a particular economic model of this kind could or should be decisive in real decision-making. Such decisions involve a whole raft of contributions and, in particular, could and should focus most strongly on specific types of possible consequences around the world.

We do, however, demonstrate in this paper that our broad assessment of the aggregate, monetary cost of inaction in the model we used is robust. There are many assumptions that could be made differently, some driving estimates downwards, others driving them upwards. In order to come to the conclusion that strong and urgent action is unwarranted, we would have to make a series of assumptions—not to mention basic omissions—that systemically and grossly bias down the cost of inaction.

Central to any aggregate assessment are ethical considerations regarding future lives and the value of additional consumption to different members of society. We recognise that the ethical discussion, which the Review helped to reignite, is extremely valuable and deserves the detailed analysis it has received. This debate should continue and is discussed further in this paper. Ethics is a subject on which reasonable people can differ, but that does not mean any old assumption will do. It is important to bring reason and evidence to the table and how that is done is of great importance. That is why ethical issues are given substantial direct attention in the Review.

We also examine criticisms of our estimates of the costs of GHG emission reductions. We find that many previous studies were too narrow in their focus. By, for example, modelling costs in the absence of induced innovation, or failing to consider the broad swathe of substitution options across countries, sectors and gases, some (but not all) previous studies overestimated the costs of mitigation.

The papers by Byatt *et al.* (2006) and Tol and Yohe (2006) that were published in the previous issue of *World Economics*, not only make many errors, they fail to bring to the table with anything like the necessary emphasis key subjects in economics, which are at the heart of the analysis of climate-change policy. In particular, these include the economics of risk, the modern public economics that focuses on imperfections in markets and limited policy tools, and the relation between economic policy-making and ethics. In some respects, although there are very important exceptions, this reflects the deficiencies of much of the preceding economic literature on climate change. These subjects underpin the analysis of the Review, although the need to communicate to a broad audience meant that they could not always be expressed with the mathematical and analytical detail required. They do, however, play an important role.[3]

Finally, we should emphasise that, whilst the first half of the Review deals with the economics of action and inaction (and it is only Chapter 6 that focuses on formal 'Integrated Assessment Modelling'), the second half of the Review deals with mitigation policy, adaptation policy and international collective action. The second half of the Review is in many ways the more important and we take the discussion further in the companion paper in this issue.

2. The cost of inaction

A basic conclusion of the Stern Review is that the costs of strong and urgent action on climate change will be less than the value of the impacts of climate change under BAU. The case for immediate action is driven by the science—by the flow-stock mechanics that determine atmospheric concentrations of GHGs. Unabated climate change poses serious environmental, economic and social risks, and the ratchet effect of the flow-stock process, particularly since carbon dioxide (CO_2) has such a long residence time in the atmosphere, makes delay costly. The case for the strength of action is also driven by the science. This section will show that, contrary

[3] Risk is crucial throughout the economics of climate change. The relations between: (i) economic activity and emissions; (ii) emissions and the stock of GHGs; (iii) the stock of GHGs and temperatures; (iv) temperatures and climate; (v) climate and impacts and so on are all stochastic. Yet the previous literature, with some exceptions, has focused on deterministic relations, thereby badly underestimating impacts (see below). Market imperfections, particularly concerning information, complicate public policy throughout, yet have generally received inadequate attention in the literature. And similarly, assumptions on intergenerational and intragenerational equity have not been afforded the necessary attention.

to some criticisms, the Stern Review assumptions on the science are unbiased and, if anything, conservative. As economists, we did not try to second-guess the science. It is the physical impacts of climate change, assessed in detail and without aggregation and monetisation (see Chapters 1, 3, 4 and 5 of the Review), that form the core of our case for action. Our basic argument is that we would readily accept a price of around 1% of GDP to avoid most of the impacts of climate change, and significantly ameliorate the risks.

2.1 Stabilisation and its flow-stock mechanics

A defining feature of climate change is the flow-stock mechanism that drives atmospheric concentrations of GHGs. This dictates that the costs of action will, under any realistic expectation of exogenous technological innovation, almost certainly rise for every additional year of inaction—in other words, delay will simply add to the costs of mitigating and preventing climate-change damages (in this respect, Tol and Yohe, 2006, appear to agree with us).

It is the stock of GHGs that correlates with increases in global temperature and consequent risks of damages. Yet it is the flow of GHGs, in the form of annual emissions, that corresponds with human activity. The problem is that GHGs stay in the atmosphere for many decades (the atmospheric residence-time of CO_2 is a century or more), so the damages that afflict the planet are the cumulative consequence of decades of activity. The concentration of GHGs in the atmosphere is now rising at a faster rate than ever before.

In the Review, we argued that the economics of climate change points to the need to stabilise the atmospheric concentration of GHGs at between 450 parts per million of CO_2 equivalent (ppm CO_2e) and 550 ppm CO_2e.[4] We are already close to 450 ppm CO_2e, and 550 ppm CO_2e gives a 50:50 chance of eventual warming above 3°C[5]—a very risky place to be (see parts I and II of the Review). Correspondingly, on an emissions path heading above 550 ppm CO_2e, the expected benefit from abating a tonne of emissions—the avoided damages and adaptation costs—is likely to be above the cost of doing so. Below 450 ppm CO_2e, the incremental

[4] See Sections 13.6 and 13.7, which explain how the different analyses of costs and damages fit together, to which Byatt *et al.* (2006) and Tol and Yohe (2006) seemed to have paid scant attention.
[5] Above the pre-industrial global mean temperature.

expected cost of abating a tonne of emissions is likely to be above the expected benefit of doing so. By our estimates, stabilising in this range will eliminate more than 90% of the cost of climate change relative to BAU. Arguments on our preferred stabilisation range are set out more formally in Chapter 13, employing findings from a wide range of modelling studies.

Unfortunately, this opportunity to stabilise the atmospheric concentration of GHGs will not wait for us. Because the stock of GHGs continues to grow, the cost of attaining a given stabilisation level increases with time. The world currently emits over 40 gigatonnes (Gt) CO_2e in GHGs each year. This is enough to raise the stock of atmospheric GHGs, which currently stands at about 425 ppm CO_2e, by about 2–2.5 ppm per year. In order to stabilise the stock of GHGs, annual emissions must be brought down towards 5 Gt CO_2e; this represents the Earth's natural capacity to absorb GHGs in any one year.

It is still possible to follow a path to stabilise at 550 ppm CO_2e. This is shown in Figure 1; emissions can peak in the next twenty years, with a smooth decline of 1%–2.5% annually thereafter. Ten or twenty years ago, a similarly smooth and affordable path might have been available for a corridor consistent with stabilising below 450 ppm CO_2e. But it is now too

Figure 1: Emission corridors to stabilisation

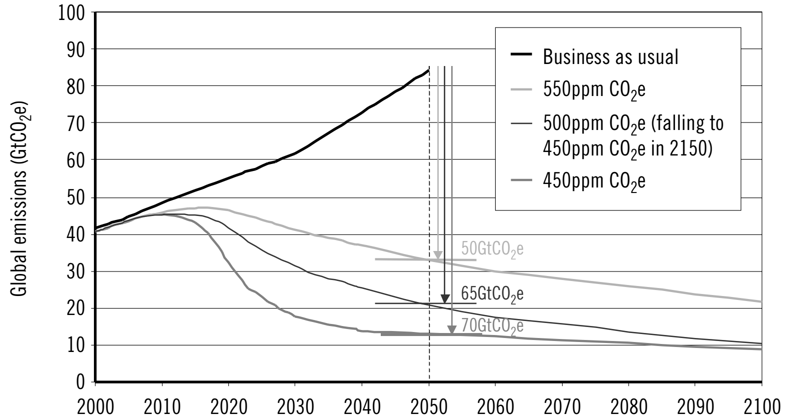

Stabilising below 450pm CO_2e would require emissions to peak by 2010 with 6–10% p.a. decline thereafter. If emissions peak in 2020, we can stabilise below 550ppm CO_2e if we achieve annual declines of 1–2.5% afterwards.

late—the kind of retrenchment required to stabilise at below 450 ppm CO_2e is likely to be extremely costly. It would require early scrapping of functional capital and the premature use of expensive technologies (whose costs will not have benefited from learning and experience). This is why 450 ppm CO_2e currently marks the lower bound of our target range. If action is delayed for another ten or twenty years, stabilisation at 550 ppm CO_2e will slip out of our reach too.

We are already committed to climate change even if we stabilise GHG emissions, so adaptation must play a role. Criticism that we have ignored or underemphasised the role of adaptation is odd (e.g. Byatt *et al.*, 2006; Mendelsohn, 2006). We devote one of the Review's six main parts to adaptation policy (Chapters 18, 19 and 20). Furthermore, our aggregate modelling includes an important role for adaptation, and in doing so explicitly, rather than implicitly, goes further than most previous studies (see Section 3).

2.2 Science assumptions consistent with received literature

Analyses of the costs of inaction and of action are based on the science that links GHG concentrations and climate. We undertook no new science (we are economists). We drew our assumptions from a broad view of the world's leading climate science. The IPCC process is perhaps the most important synthesis of this. Its *Assessment Reports*, published around once every six years, are widely recognised to be the most comprehensive of their kind. The process is not perfect, but we find no cause to throw its principal conclusions into question. Indeed, we followed IPCC in its caution, which leads it to leave out from formal estimation and prediction a number of risks thought to be possible, but on which insufficient research has been done to provide strong probabilistic estimates. That caution, which is understandable, leads to possible underestimation of climate impacts, as we illustrate in the Review (see, e.g., Box 6.2).

Nevertheless, we knew the science had moved on since the 2001 *Third Assessment Report* (TAR) and, in making our projections, we took on board the latest evidence (in particular, Murphy *et al.*, 2004, and Meinshausen, 2006). As well as being more up-to-date, these include a more sophisticated treatment of the probabilities of temperature rise, which is a crucial input to the economics of risk. We have received criticism for departing from the 2001 TAR assessment of the links between concentrations and temperature increase, which was based mostly on late-1990s science. The fact that, subsequent to publication of the Review, the IPCC *Fourth Assessment Report* (AR4: 2007) updated its earlier results and outlined conclusions similar to ours was no surprise, as we drew on the same body of evidence and consulted the same leading experts.

As time has progressed, scientists have acquired a better understanding of the likely non-linearities, feedbacks and convexities in the geo-climatic system. Very importantly, they are also developing probabilistic assessments of the consequences of climate change. These are the key reasons why estimates of the total risks have tended to increase with time. The upper bound of the IPCC's range of predicted global mean temperatures has been driven higher than in the TAR by the inclusion of positive feedbacks from the carbon cycle, such as a weakened uptake of natural carbon by forests, vegetation and soils, as areas become warmer and drier. The findings of the IPCC are now strongly consistent with the scientific evidence summarised in the Stern Review for projected temperature rise, patterns of precipitation, and the incidence of extreme events. A detailed comparison of our results with those of AR4 can be found in Paper A, 'The case for action to reduce the risks of climate change', on our website.[6]

We have further been accused of overstating the damages of any given amount of global warming (Byatt *et al.*, 2006; Tol and Yohe, 2006), either because we have used models that are clear outliers, or because we have imposed biased assumptions. In Section 3, we will show that the assumptions used are, if anything, conservative, and that our aggregate model properties are sensible and consistent with the leading alternative studies.

2.3 Understanding the physical impacts of climate change

We built our assessment of the costs of inaction (or the benefits of climate-change mitigation) around a disaggregated analysis of physical impacts around the world. Chapters 1, 3, 4 and 5 focused on changes in temperature, the water cycle, sea-level rise, extreme events, and biological, chemical and physical processes of global importance (such as plant and soil 'sinks' for carbon, and ocean acidity, with its consequences for marine life). The key question that arises is whether paying as little as 1% of GDP over much of this century (and possibly beyond, depending on technical

[6] www.sternreview.org.uk

advance) is worthwhile to reduce the risks and uncertainties described. We argue that a balanced judgement would conclude that it is.

Climate change threatens the basic elements of life for people around the world—access to water, food, health, and the use of land and the environment. The industrialised world is not immune to climate change, despite the fact that less of its output is produced by primary, climate-sensitive sectors such as agriculture and forestry. To claim otherwise (see Byatt *et al.*, 2006) is to ignore the clear links between climate and secondary and tertiary economic sectors such as housing, transport and other infrastructure, insurance and financial services. Many of the potential physical changes such as extreme weather and sea-level rise would have economy-wide effects, not to mention the many ways in which climate affects health, and ecosystems contribute to social welfare. Finally, the risks associated with large BAU temperature increases of 5°C and upwards (see Chapters 1, 3, 4, and 5 of the Review) have the potential to disrupt life on the planet at a level that is hard to imagine and prepare for.

At 1°C–2°C temperature rise, there will be some winners and some losers. Longer growing seasons in northern latitudes, and reduced mortality from winter cold snaps, will create economic gains in some areas and opportunities for new activities, including in the agriculture, energy and tourism sectors. But even at low levels of warming, there are already significant impacts on vulnerable communities, for instance in indigenous Arctic communities and low-lying Pacific islands. Water stress in semi-arid regions and extreme events, particularly in the sub-tropics, are increasingly being linked to global warming. But unabated climate change in the long run is a different matter altogether. It requires that we seriously evaluate the consequences of 5°C–6°C warming or more, and where sea levels are rising further and faster than under more moderate scenarios. 5°C is the difference between now and the last Ice Age. A further 5°C would transform the Earth's physical geography, putting its human geography under severe pressure and likely to lead to large-scale and very disruptive movements of population.

The paleoclimatic record suggests that climate change has not taken place in a linear, gradual way. Significant changes to the global climate have occurred within only a couple of decades. Regional events that could bring severe disruption with little advance warning include a strengthened El Niño event or widespread forest fires in Siberia or the Amazon. These could trigger an abrupt failure in monsoon rains and a significant fall in agricultural yields in key areas of Asia, Australia or Latin America, with implications for the global trade in commodities such as wheat and soya, as well as risks of human misery, social instability, and migration in densely populated regions of the world. These events do not depend on crossing a known critical threshold of global average temperature, but become ever more likely as global temperatures rise. Furthermore, there are a number of basic biological and physical principles indicating that impacts in many sectors will become disproportionately more severe with increasing warming. Thus, to give just one example, recent estimates suggest a ninth-power relationship between hurricane wind-speed and damages (Nordhaus, 2006a). Hurricane intensity is expected to rise as a result of warming (although changes to their location and numbers remain less certain).

At this detailed level of analysis, the impacts of existing climate and weather patterns are a significant source of evidence, for example on the consequences of extreme weather events. Thus the criticism that the Review 'gives too little attention to actual observation and evidence, as distinct from the results of model-based exercises' (Byatt *et al*., 2006, p. 224) is unwarranted. Nevertheless, BAU climate change is very likely to take us into uncharted climatic territory, to conditions beyond human experience, so there can ultimately be no substitute for forecasting into the future with the tools that we have. The economics of risk requires that we look beyond what we think we know from past evidence to ask what might happen in the future.

3. Integrated assessment models and the aggregate cost of inaction

Despite basing our case for action on detailed, disaggregated projections of climate-change impacts, disproportionate attention has been paid by some economists to the aggregate studies outlined in a single chapter of the Review, Chapter 6. While understandably of great interest to economists, the economics of climate change does not rest solely on estimates of the aggregate, monetary value of climate-change impacts: thus many critiques focusing on Chapter 6 are fundamentally misplaced in their emphasis. Aggregate estimates are highly simplified and suppress much, or most, of what is interesting and troubling about climate change. Nevertheless, they do help explore some (but not all) aspects of the logical structure behind valuing the impacts and risks of climate change, and require an explicit statement of ethical considerations used in such an analysis.

This section outlines how the Stern Review's aggregate estimates were derived, focusing particular attention on the most important assumptions and judgements that must be made, *inter alia* on discounting and social welfare valuation, the treatment of risk and uncertainty, and adaptation. This allows us to set out a framework for sensitivity analysis and an assessment of the robustness of our basic conclusion that the costs of strong and urgent action are much less than the costs of inaction. But first we clarify the relationship between the estimates of the Integrated Assessment Model (IAM) we used, PAGE2002 (fully described in Hope, 2006), and its principal peers.

3.1 PAGE2002 and other models

We chose the PAGE2002 IAM, in large part because it is stochastic, embodying Monte-Carlo procedures to estimate probabilities, and calibrated to reflect the range of disagreement and uncertainty in the underlying scientific and economic literatures. For each of the model's some 31 inputs that are stochastic, a probability distribution is drawn from the corresponding range of estimates in the applied literature. So we allow for a wide range of 'stories', from optimistic to pessimistic. Thus we can demonstrate that the costs of climate change estimated by PAGE2002, as a dynamic function of global mean temperature *before* social welfare valuation, lie at the centre of a range of studies: it is not an outlier. Figure 2 plots the undiscounted cost of climate change as a function of global mean temperature for PAGE2002 (the mean estimate in each time period), in comparison with three other leading IAMs.[7] Costs are expressed relative to global output projected in a world without climate change (i.e. with the impacts of climate change simply 'switched off').

Crucially, the four IAMs under examination here are not equal in their coverage of potential climate impacts. Mendelsohn *et al.* (1998) choose to omit all but impacts on a handful of market sectors of the economy.

[7] The estimates of Nordhaus and Boyer (2000)—the 'Nordhaus' estimates—and Tol (2002) are reported twice, for different methods of weighting impacts worldwide.

Figure 2: Comparing the dynamic costs of climate change, as a function of global mean temperature, estimated by four leading IAMs

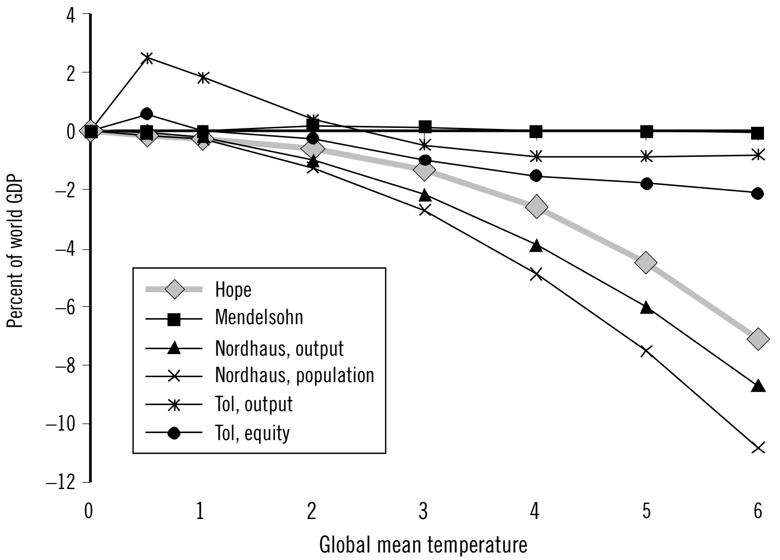

Source: Adapted from figure 6.2, p. 166, of Stern (2006), with original data from Smith *et al.* (2001)

By inference, there are no direct, welfare-equivalent impacts on human health and ecosystems whatsoever in their study. There is a rationale for this—the monetary valuation of such 'non-market' impacts is more speculative and uncertain. But it is essential to interpret the results of such studies for what they are: estimates of a narrow subset of possible climate impacts. Tol (2002) includes a much wider range of impacts, but restricts himself to gradual climate change. Nordhaus and Boyer (2000), on the other hand, in addition to estimating both market and non-market impacts, also include an estimate of the risk of abrupt, large-scale and discontinuous changes to the climate system, for example a regional or global shut-down of the Thermohaline Circulation, or rapid melting of the Greenland and West Antarctic Ice Sheets. Recent scientific evidence is increasingly pointing to such risks. PAGE2002—the 'Hope' estimates—also includes estimates of market impacts, non-market impacts and the risk of a large-scale discontinuity, making it as comprehensive as any of its peers.

Yet we must reaffirm that *all* of the models make potentially important omissions: none incorporates estimates of all the impacts of climate change considered possible (see Downing *et al.*, 2005, and Figure 6.3 in the Review). For example, none explicitly takes into account interactions between impacts in different sectors, a process which has been identified in multi-sectoral (e.g. general equilibrium) economic modelling studies (see Jorgensen *et al.*, 2004). Furthermore, none explicitly takes into account so-called 'socially contingent' impacts, which are large-scale, 'second-round' socio-economic responses to climate change like conflict and migration (Byatt *et al.*, 2006, are mistaken in their suggestion that we 'ramped up' estimates by adding damages due to extreme weather, 'social and political instability' and 'knock-on effects').

3.2 Stern Review aggregate estimates

Critiques of the Review have portrayed our estimates as a clear outlier in relation to the previous literature. To begin with, this is misleading, because it ignores the increasing number of sensitivity analyses published in recent years, which have explored uncertainties in scientific and economic research. These represent a partial attempt to quantify the great uncertainties in estimating climate-change impacts, and also produce high estimates of the cost of BAU climate change (e.g. Ackerman and Finlayson, 2007; Azar and Sterner, 1996; Azar and Lindgren, 2003; Ceronsky *et al.*, 2005; Roughgarden and Schneider, 1999), sometimes far higher than those of the Review. We shall have cause to return to the treatment of uncertainty below, because our approach goes further than sensitivity analysis, towards a systematic valuation of all quantifiable risks. In addition, we have just demonstrated that, contrary to some critiques (e.g. Byatt *et al.*, 2006), PAGE2002 is not an outlier, neither in the impact sectors it includes, nor in the simple, undiscounted estimates it produces. However, we were not surprised that our estimates are higher than many of their counterparts in the previous literature. It is our view that such estimates downplayed the risks of future climate change and their present value. Here we outline why.

3.3 Discounting and intergenerational equity

Much discussion of the Review has focused on discounting, and rightly so. It is clear, and was very strongly emphasised in the Review, that a very low weight on the future will simply downplay the risks that occur 50, 100 or 150 years from now (see Chapters 2 and 6). An examination of ethics and value judgements cannot be avoided when considering the weight attached to the impacts of global warming, many of which are long term. (A detailed account of our approach to discounting can be found in Paper B, 'Value judgements, welfare weights and discount rates: issues and evidence', on our website, and in Chapter 2 and its Appendix in the Review.)

Questions postulated in terms of '*the* discount rate' are often badly formulated. As we explained in Chapter 2 and its Appendix, when the consumption consequences of policy choices could be non-marginal, the marginal methods embodied in a discount rate are misplaced. Climate change has the potential to cause permanent, large-scale losses in consumption, so to evaluate them in terms of a perturbation around a path is misleading. That is why we worked directly in terms of the evaluation of an expected utility integral.

Nevertheless, discussion of discount rates can be useful in understanding some of the issues that shape our perspective of the relative weights to be ascribed to current and future benefits and costs (such as the relative roles of growth and of pure disregard for the future). In this context, we set out in detail in the Appendix to Chapter 2 of the Review one particular approach to intergenerational ethics. The discount factor (the relative weight on an investment in consumption at time t, relative to now) in this context is given by $u'(c)e^{-\delta t}$, where $u'(c)$ is the marginal utility of consumption per head at time t and δ is a 'pure time discount rate' (also known as the utility discount rate or pure rate of time preference) describing a lower weight on the future, *simply because* it is in the future. This is fairly standard in the literature. The consumption discount rate r (not to be confused, as some have, with the pure time discount rate) is the rate of fall of the discount factor:[8]

$$r \equiv \eta g + \delta \quad (1)$$

where η is the elasticity of marginal utility of consumption and g is per-capita consumption growth. Equation (1) applies to small changes

[8] This discussion (and see Appendix to Chapter 2) explains why discounting associated with public infrastructural projects—for example as described in the HM Treasury Green Book (HM Treasury, 2003)—cannot apply to irreversible, non-marginal decisions with impacts in the far-off future. The choice of what action to take on climate change will itself have a bearing on the appropriate discount rate that needs to be used.

(increments) around a particular path. It does not apply to non-marginal changes and, where applicable, it depends on the particular path chosen for study and will vary over time.

The most straightforward and defensible interpretation (as argued in the Review) of δ is the probability of existence of the world. In the Review, we took as our base case $\delta = 0.1\%$ per year, which gives roughly a one-in-ten chance of the planet not seeing out this century.[9] For much of the Review's discussion, we set η equal to unity, although we considered sensitivity in the Postscript to the Review. That choice of η is fairly standard, but the role of η is discussed below. Across the infinite horizon of our analysis in Chapter 6, g is on average around 1.3% in a world without climate change, giving an average consumption or social discount rate across the entire period of 1.4% (being lower where the impacts of climate change depress consumption growth).

In Table 1 below, we briefly summarise the sensitivity of our estimate of the total cost of BAU climate change to variations in δ and η. If both are increased, the consumption discount rate increases and this reduces our estimates of the cost of BAU climate change significantly.[10] This is unsurprising. It says little more than 'if you don't care much about the future, you won't care much about climate change'. Some critics of this part of the Review favour higher consumption discount rates (e.g. Dasgupta, 2006; Nordhaus, 2006b; Tol and Yohe, 2006). The whole approach, while fairly standard, does depend on value judgements about intergenerational equity and, within that, choices surrounding η and δ. Let us consider the issues briefly.

Much of the substance of the criticism surrounds whether it is appropriate to follow a 'prescriptive' approach to discounting, or a 'descriptive' approach that seeks to present the consumption discount rate as a derivation of the revealed preferences of the current generation. The latter would likely favour higher values of δ. The implications for η of a descriptive approach are ambiguous (see Paper B on our website), because in the fairly standard framework for social welfare valuation that we followed

[9] This may sound pessimistic, and those who feel it is will discount by a smaller amount. Others might opt for greater existential discounting: for example Martin Rees, who, in *Our Final Century* (Rees, 2003), argues that there is only a 50:50 chance of making it to the end of the century; this is equivalent to discounting at an annual rate of 0.7%.

[10] Further sensitivity analysis is reported in the Postscript to the Review, and its Technical Annex, as well as Paper A on our website.

(utility functions with constant η), the rate of intergenerational inequality aversion is equal to the rate of intragenerational inequality aversion, and both are equal to the rate of (relative) risk aversion. Evidence from each of these three dimensions can be contradictory (see, e.g., Paper B, Quiggin, 2006, and references in the Review). Disentangling these three issues is an important topic for future research.

The valuation of long-run climate risks cannot in any case be driven by a simple, descriptive approach to discounting. There are severe limitations to an approach, which relies on actual market behaviour to 'reveal ethics' (Hepburn, 2006) in the manner required for the long-run issues raised by climate change. Examination of rates of discount associated with projects of a shorter time-scale, or which reflect aggregated private decisions made over individuals' lifetimes, are inappropriate to describe a sensible ethical treatment of risks and uncertainties caused by irreversible actions today, which befall generations more than 100 years in the future.

The issues are particularly striking in relation to δ, or the rate of pure time preference. Pure time discounting is, in effect, discrimination by date of birth. For example, $\delta = 3\%$ would give individuals existing at the end of this century roughly one tenth of the ethical weight of the current generation, irrespective of their relative income, and individuals existing in 2200 would receive only one hundredth of the same weight. For $\delta = 2\%$, someone born in 2008 would have around half the ethical weight of someone born in 1973. It is hard to see why the impatience of some people in their own lives for jam today rather than jam tomorrow justifies the extreme ethical discrimination by date of birth in the examples just described. There is no market place reflecting an answer to the direct question 'how should this generation, deciding together, act to allocate resources and protect the environment for the next generation?' And there is no market that can allow us easily to distinguish pure-time discounting from discounting. In any case, evidence from the market place itself is very mixed. In the UK, many have recently bought index-linked 50-year government bonds with real interest rates of 0.4% per annum.

It has been argued that a simple application of $\delta = 0.1\%$ per year and $\eta = 1$ could imply very high optimum saving rates (see, e.g., Dasgupta, 2006; and Nordhaus, 2006b). In cases such as this, the implication just described is highly sensitive to model structure. The arguments of Dasgupta and Nordhaus depend very heavily on the assumption that there is no technical progress. With technical progress, much less saving is required for η at unity to be consistent with growth along an optimal path.[11]

Nordhaus argues for higher δ and Dasgupta for higher η, based on this kind of argument. We have argued strongly that the ethical case for higher δ seems weak. That for η concerns distributional judgements. On this there can be many views. But if one pushes the case for $\eta = 2$, say, it is important to recognise that one should be in favour of strong redistributional policies within a generation. If A has 5 times the resources of B then, with $\eta = 2$, a marginal unit of resource for B is worth 25 times that of A (since $u'(c)$ is the reciprocal of the square of c). Thus we could lose 95% of a unit transferred and would still regard the transfer as justified.

Some have argued that with $\eta = 1$ and $\delta = 0.1\%$ p.a., a (very long-run) consumption discount rate of 1.4% is implausibly low, since there will be a social return on investment much higher than this. They argue that we would be better off investing elsewhere than in avoiding climate change. This argument is predicated on the apparent knowledge of what very long-run rates of return might be. It is hard to know why we should be confident that social rates of return would be, say, 3% or 4% into the future. In particular, if there are strong climate change externalities, then social rates of return on investment may be much lower than the observed private returns on capital over the last century, on which suggestions of a benchmark of 3% or 4% appear to be based. The argument also neglects that such a strategy, which proposes action to prevent climate change, will face much higher costs, when it tries, much later, to tackle climate change. Indeed, many impacts will be irreversible. Further, a higher social discount rate can come from a higher rate of growth g (see equation (1)). Increasing g in these models not only increases discount rates, however, it advances the pace of atmospheric GHG accumulation and, in this respect, makes the request for action more urgent.

Whilst we have found the criticisms of the Review based on η and δ to be less than convincing, we must recognise that no discussion of the appropriateness of particular value judgements can be decisive. It is for this reason that we have conducted several sensitivity analyses. Societies

[11] Brad DeLong demonstrated this on his 'blog' website. The importance of model structure was also explained in the Review where this type of issue was discussed directly—see Appendix to Chapter 2, p. 54. Nevertheless, for $\eta = 1$ and $\delta = 0.1\%$, one of the conditions for the convergence of utility integrals is only just satisfied (see p. 58). This does lead to strong weight on future damages. Whilst many might regard this as appropriate, others might not. That is why we provided sensitivity analysis—see pp. 658–71 of the Review.

around the world should embark on a serious and sustained discussion of the ethical dimensions of climate policy. Indeed, perhaps the popularisation of issues around discounting in the media in recent months indicates that they have.

3.4 Treatment of risk and uncertainty

All the links in the chain between GHG emissions and the economic impacts of climate change—each of which needs to be parameterised in an IAM—are subject to uncertainty. Yet aggregate studies have not always chosen to tackle this uncertainty directly. The simplest modelling strategy in the literature is deterministic, whereby a 'best guess' is made for each parameter. This is still very common. Most IAMs have also been set up at one time or another to run a Monte-Carlo procedure. This enables climate impacts to be modelled probabilistically (e.g. Hope, 2006; Mastrandrea and Schneider, 2004; Plambeck *et al.*, 1997; Roughgarden and Schneider, 1999). Yet very few of these studies extend to a full application of expected-utility analysis (exceptions are Tol, 1999 and 2003), which enables a valuation of relative climate risks against an overall criterion. While expected-utility analysis is often used to investigate issues around learning and the resolution of uncertainty over time, it is most surprising to us that it has not become the standard method of social welfare valuation in this more simple exercise—estimating the cost of inaction under BAU.

In Table 1 below, we demonstrate the proposition that inadequate treatment of uncertainty leads to estimates of the cost of BAU climate change that are misleadingly small. We compare the estimate that PAGE2002 produces using expected-utility analysis to that produced when we adopt a deterministic, 'best guess' approach, setting each model parameter to its mode (most likely) value. The difference is large—of the same magnitude as those produced by very different assumptions about discounting. Recent scientific evidence increases the imperative to adopt an expected-utility approach, since it has raised the possibility of very large impacts and as a result has effectively increased the confidence interval around the future consequences of climate change. Further, it now gives much more detailed estimates of probabilities, allowing the application of the economics of risk.

Despite this, it is very likely that even PAGE2002 underestimates the uncertainty around many of its parameters, because of limitations in the data it must rely on. In many cases, the probability distributions that are estimated for its parameters are based on a range of underlying studies, which themselves give 'best guesses'. As such, PAGE2002 can encapsulate uncertainty between the best guesses of other models, but it is unlikely to adequately capture uncertainty within these models themselves. It is in this context that Tol and Yohe's (2006) accusation of double-counting risk looks entirely misguided (see also the appendix to this paper). The situation would better be described as 'under-counting' risk. It is not only that risks on which there is some awareness are excluded. There is also strong cause to suspect that rising stocks of atmospheric GHGs carry with them the risk of consequences yet to be discovered. Since 'nasty surprises' are widely expected to be more likely than 'nice surprises', this is a firm justification for strong GHG emission reductions—for precaution. Going still further, we are likely to require in this context the theories of choice under risk and uncertainty, which distinguish between risk (known probabilities) and uncertainty (unknown probabilities) in the Knightian manner—see p. 38 of the Review.

3.5 Adaptation

The Review has been criticised for failing to adequately accommodate the scope for adaptation. This seems odd to us, since one of the Review's six main parts was in fact devoted to adaptation (Chapters 18, 19 and 20). We emphasised the importance of adaptation as a policy response, particularly to the unavoidable impacts of climate change in the immediate future. Yet evaluating the role of adaptation in the medium to long run requires careful consideration of what it can achieve, in the face of rapidly increasing BAU warming, as well as how much it will cost to adapt.

Adaptation is also a key feature of our aggregate modelling of climate impacts. In fact, PAGE2002 goes beyond most existing IAMs in making the contribution of adaptation explicit. This allows a more transparent investigation of the effect of different assumptions about adaptation. In most other IAMs (e.g. Nordhaus and Boyer's DICE-RICE-99 models, described in their book from 2000), adaptation is an implicit, inseparable component of the function that describes damages at different levels of warming. The cost of adaptation is a cost of climate change and, symmetrically, one of the benefits of climate-change mitigation is a reduction in adaptation costs. It is not clear to us whether other IAMs, for which adaptation is implicit, count saved adaptation costs as a benefit of mitigation.

PAGE2002 assumes that, in industrialised parts of the world, adaptation will reduce the impacts of climate change on market sectors of the economy such as agriculture by 90%, at all levels of warming. In lower-income regions—Africa, India and Southeast Asia, and Latin America—adaptation reduces market impacts by 50%, irrespective of warming. Worldwide, non-market impacts, primarily on human health and ecosystems, are reduced by 25% in value terms through adaptation, again at all levels of warming. This lower non-market adaptive capacity partly reflects the insuperable challenges facing natural systems in adapting to rapid shifts in habitat conditions. Adaptation is highly cost-effective in PAGE2002, coming at a fraction of the cost of avoided impacts as global mean temperatures rise to higher levels.

We do not consider these to be pessimistic assumptions about the scope for adaptation. On the contrary, they may be optimistic. The impacts of climate change are expected to rise faster than global mean temperature, so that, even if higher incomes and advanced technologies increase capacity to adapt, BAU climate change will just increase what we have to adapt to. As we have said, 5ºC warming will be transformational. Whether we will still be able to reduce impacts on market sectors by 90% is very much open to question. Even if the capacity is in place, there is a strong argument that the cost of adaptation will also rise faster than temperatures; that is, adaptation costs are convex. For example, at small temperature rises, agricultural techniques can be altered, buildings modified and levies extended and enhanced. At higher temperatures, these costs are likely to rise more than proportionately, as coastal areas are evacuated, desertified farms are abandoned, and commercial and residential infrastructure is re-built. In PAGE2002, however, adaptation costs are a concave function of global mean temperature. Future research should reappraise this relationship.

For now, we include in Table 1 a direct response to the criticism of Tol and Yohe (2006). In their critique, they focus on our assumption of 'vulnerability that is…constant over very long periods of time' (p. 233). First, we must clarify what they mean, because it is ambiguous. In PAGE2002, vulnerability is constant in *relative* terms. The percentage impacts of climate change that we can adapt to is constant, but since the impacts of climate change rise under BAU, so does the absolute contribution of adaptation. Nevertheless, they appear to object to the idea that Africa, India and Southeast Asia, and Latin America are always less able to adapt than richer parts of the world. So in Table 1 we compare our base case with an alternative scenario, in which the vulnerability of these regions instantly falls to the lower level of the European Union, the model's base region, in 2100. We do not find that our estimates are particularly sensitive to this change.

3.6 Modelling high-damage scenarios

As we argued in the Review, there is strong cause to believe that the impacts of climate change will increase faster than temperature, as emissions drive warming beyond the range of 2°C–3°C above pre-industrial. The primary focus of most previous studies on this level of warming has led them to ignore many of the risks posed by climate change. The Review went further than most previous studies in including three pathways through which the costs of climate change can escalate as warming rises. Recent evidence has pointed ever more strongly to the risks of escalating damages and, in so doing, has indicated that the evidence base used in most IAM studies, predominantly from the 1990s (Warren *et al.*, 2006), is becoming out-of-date. As always, we model these three pathways probabilistically, thus allowing us to account for the significant uncertainty around them.

The first pathway comprises rapid, large-scale impacts of gradual climate change. In the most aggregated IAMs like PAGE2002, the impacts of gradual climate change are a function of global mean temperature. Generally, this function is calibrated first through an estimate of overall impacts at 2.5°C or 3°C warming, and second through an estimate of the functional form. A very simple specification of the damage function, in this case with respect to 2.5°C warming, is as follows:

$$d(t) = \beta \left(\frac{T(t)}{2.5} \right)^{\gamma} \qquad (2)$$

where T is warming at time t, in terms of global mean temperature above pre-industrial, $d(t)$ is the economic damage caused by climate change expressed as a fraction of consumption or income, β is the consumption loss accompanying 2.5°C warming and γ is the damage function exponent (details can be found in the Appendix to the Postscript to the Review).

While IAM studies assume small aggregate damages for initial warming of 2.5°C, they differ in their assumptions about γ, and this drives very big differences in estimates of damage at higher global mean temperatures.

PAGE2002 draws γ from a triangular probability distribution with a minimum of 1, mode of 1.3 and maximum of 3 (giving a mean of about 1.8). This captures the range of assumptions made in other studies (e.g. Nordhaus, 1994; Nordhaus and Boyer, 2000; Peck and Teisberg, 1992; Roughgarden and Schneider, 1999) and gives a small likelihood of dangerous economic impacts. In Table 1, we emphasise how important assumptions about this particular parameter are. We compare the base case distribution of γ with a scenario, in which γ is fixed equal to its maximum value of 3. The sensitivity of the cost of climate change due to γ is very high indeed.

The second pathway comprises the large-scale release of GHGs from sinks, which constitutes a positive natural feedback to global warming by accelerating the overall atmospheric build-up of GHGs. These positive natural feedbacks are uncertain and as such have received virtually no attention in IAM studies to date (but see Ceronsky *et al.*, 2005). But scientific evidence is pointing more and more strongly to the risks. The Review thus goes further than almost all previous studies in specifying a high-climate scenario to take account of recent quantitative modelling of positive natural feedbacks (e.g. Friedlingstein *et al.*, 2006; Gedney *et al.*, 2004). We present this beside our baseline-climate scenario, which is calibrated on the overall conclusions of the 2001 IPCC TAR, but also includes small positive feedbacks (see Hope, 2006). Contrary to how the high-climate scenario has been portrayed by some critics, it is, in fact, fairly conservative, because it deliberately omits some of the positive feedbacks currently under examination in the science (e.g. the large-scale release of gas hydrates from the oceans). The combined effect of the feedbacks is to increase mean warming in 2100 by only 0.4°C, which is slightly less than the effect projected in the IPCC AR4 (IPCC, 2007). Table 1 compares the difference in our mean estimate of the total cost of BAU climate change between the baseline- and high-climate scenarios.

Pathway three comprises abrupt, large-scale and discontinuous changes to the climate system with direct economic impacts, such as a regional or global shutdown of the Thermohaline Circulation, and rapid melting of the Greenland and West Antarctic Ice Sheets. These non-linear changes are difficult to predict, but are plausible given what is known about the chaotic nature of the climate system and past climate changes. So far, only two IAMs have built the possibility of these kinds of climatic changes into their core model structure. The pioneering estimates were included in

DICE-RICE-99 (Nordhaus and Boyer, 2000). PAGE2002 also includes the risk of a climate discontinuity. Table 1 presents our mean estimate of the total cost of BAU climate change with and without this risk.

3.7 Sensitivities—what drives the results?

In Table 1, we summarise the sensitivity of the Review estimates to the four key issues discussed above:

1. Ethics and discounting;
2. Treatment of risk and uncertainty;
3. Adaptation;
4. Modelling high-damage scenarios.

In each case, we have carried out formal sensitivity analysis, using PAGE2002. More detail is given in the Technical Annex to Paper B, on our website. The base case in Table 1, from which deviations are reported, is our 'central' modelling case. 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 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.

For each of the modelling issues we have examined, we summarise the effect on our mean estimate of the total cost of BAU climate change when one parameter is varied, holding everything else constant.[12] 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). 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).

The first two rows of Table 1 focus on ethical aspects or value judgements and the next five on model structure. Of course, the overall implications of any change in one assumption depend on given assumptions, both ethical and structural, throughout the model.

[12] Note that the effect of varying several parameters at once will not be additive.

Table 1: 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

Table 1 shows that alternative assumptions on these dimensions produce a wide range of possible outcomes, some lower than our central estimate, some much higher. At the lower end of the range, the cost of inaction is low and only very moderate GHG emission reductions are warranted. But, in order to reduce our ambitions to those implicitly advocated by Tol and Yohe (2006), or even further to those advocated by Byatt *et al.* (2006), we would have to make value judgements, assumptions about the future, as well as basic omissions to the model and the modelling approach, which systematically bias down the cost of inaction. At the upper end of the range, the cost of inaction is far in excess of our summary estimates. Stringent and immediate emission reductions, beyond those advocated by the Stern Review, would be justified in this case.

In addition to the issues we have discussed in detail, there are many other assumptions and judgements that will affect results. We summarise a selection of these in Table 2, along with a back-of-the-envelope calculation of how they will affect our estimates. Again, the first two rows focus primarily on ethics or value judgements (while population growth is not an ethical issue in the strict sense, it affects the valuation of impacts via its position in the overall social welfare function—see Box 6.3 of the Review), while the remaining five focus on structural aspects. Box 1 provides more by way of explanation. Again, a wide range of possible estimates emerges, but in most cases these variations in assumptions would increase our estimates.

Overall, from this discussion of sensitivity, there seems little justification in changing our broad view that the cost of BAU climate change, in terms of a loss in present global mean consumption per capita, is well in excess of the cost of stabilising GHG emissions in the range of 450–550 ppm CO_2e. It is very clear that IAMs produce results that are sensitive to assumptions. That is one reason why they should be used with caution and why we laid strong emphasis on the disaggregated approaches of Chapters 1, 3, 4, and 5. Nevertheless, it is clear that in many respects we were cautious about, or omitted, many aspects of the modelling structure, which would have raised damages. There is no justification in the claim that we systematically chose assumptions which would give high damages. In our view, the range of damage estimates presented in the Review is in the right 'ball park'. We would have to make a series of quite extreme assumptions, in order to overturn the claim that the costs of inaction are more than the costs of action.

Table 2: 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 modelling 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	Modelling 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 stabilise	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	++

BOX 1
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.* 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

* Our regional growth rates were converted from market exchange rates to purchasing power parities.

about by such a high stock of atmospheric GHGs, dominate over the longer term.**

Terminal conditions
In other words, the length of the modelling horizon and what is assumed to occur thereafter. The PAGE2002 modelling 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 stabilise and the stock of GHGs to rise very slowly. The longer the modelling 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 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.

** 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).

4. The costs of mitigation

Having assessed the impacts of BAU climate change in Parts I and II of the Review, in both a disaggregated and aggregated way, the Review went on to analyse, in Part III, the costs through time of keeping GHG emissions in line with a plausible stabilisation corridor—keeping atmospheric concentrations below 550 ppm CO_2e, in order to substantially reduce the risks of the highest increases in temperature. In the important closing chapter (13) of Part III, we also provided an analysis of how to combine estimates of costs of action and of inaction—see section 5 below.

Our assessment was based, first, on the latest literature covering model estimates of the cost of mitigation. Our central estimate of an annual cost of 1% of GDP to follow a 550 ppm CO_2e stabilisation pathway applies through to the end of this century (with a range of +/–3 percentage points by 2050). This is based on a very wide range of international model estimates and does not stem from the use of one particular approach.[13] In thinking about costs, we made the simple assumption that the 1% cost continues forever, although it is quite likely that over the long run it would come down sharply.

The Stern Review also commissioned a simple and transparent cost-assessment exercise, designed to complement the wide-ranging model estimates. In order to assess the likely costs of mitigation in a world where behavioural change was limited—a very conservative assumption—a probabilistic projection of the evolution of low-carbon technologies and of fossil-fuel prices was used (Anderson, 2006).

This study gave results in a similar range to the more complex behavioural modelling exercises and showed that under a feasible technology mix, substituting carbon-intensive energy generation and transportation with low-carbon technologies to meet an emissions corridor consistent with stabilisation at 550 ppm CO_2e[14] could be attained with a mean cost of approximately 1% of GDP by mid-century. The uncertainty around this mean amounted to around +/–3 percentage points of GDP, reflecting in

[13] The following studies were investigated: Stanford University's Energy Modelling Forum (EMF-16, EMF-19 and supporting multi-gas estimates for EMF-21); the meta-analysis study by Fischer and Morgenstern (2005); the IEA accelerated technology scenarios; the IPCC TAR survey of modelling results; the Innovation Modelling Comparison Project (IMCP); the draft US CCSP Synthesis and Assessment of 'Scenarios of Greenhouse-Gas Emissions and Atmospheric Concentrations and Review of Integrated Scenario Development and Application' (June 2006).

[14] Relative to a baseline path consistent with the IPCC SRES A2 path used to estimate the damages.

particular uncertainty about technological innovation and the evolution of fossil-fuel costs. Unlike the behavioural models, whose results are driven by detailed assumptions and parameters, this approach offered a very simple and transparent way of making a first approximation of the likely cost of one route to decarbonising economic activity. An approach that embodies cost-minimisation and formal optimisation would give a cost no higher than this, if it includes this approach as one possible option.

It should be noted that marginal costs are not falling through time in the Anderson study, as suggested by Tol and Yohe (2006): they are rising. Average costs are falling, and there seems to have been some confusion in their argument between average and marginal costs. Understanding the difference is crucial. Box 9.6 of the Review points out that, although marginal costs are likely to rise through time, in line with a rising marginal damage cost ('social' cost) of GHG emissions, average costs may rise, fall, or stay the same, depending on the rate of technological progress. This is fully in line with the academic literature on the evolution of costs.

There has been some debate about the estimates presented in the Review of the costs of mitigation and it is important to try to clarify some simple misunderstandings.

First, this estimate is not materially affected by the choice of discounting assumptions, as the mean flow cost-estimate of 1% of GDP is relatively constant through the medium term.[15] It is designed to be consistent with our approach to expressing likely climate-change impacts, using 'balanced growth equivalent' (BGE) paths. Both approaches to measuring impacts and mitigation costs express costs as a percentage of real consumption or income and are directly comparable. The simplest way to think of the cost is a one-off, 1% increase in a cost or price index. It applies to both a flow of consumption and to a flow of income that generates the consumption.

Second, a more interesting question concerns the validity of attempts to compare near-term costs with long-term impacts, and the issue of whether 1% of GDP represents a trivial investment. There is no inconsistency in comparing costs over the next 100 years with impacts occurring in the more distant future.[16] As we argued, the decision on whether to act now

[15] There is a range of projections of mitigation costs beyond 2050. By 2100, some fall, reflecting greater efficiencies from induced technological innovation, while others rise sharply, reflecting the greater uncertainty over the costs of seeking out successive new mitigation options.
[16] In any case, the evidence suggests that a 1% annual cost of mitigation is a reasonable estimate of the flow of costs into the long term.

hinges on the question of irreversible outcomes and risks. Decisions taken today will have potentially large and irreversible consequences in terms of climate-change impacts; this is not true to the same extent of mitigation costs.[17] Moreover, policy-makers can keep cost estimates under review and revise policy in the light of new information. By contrast, the impacts of climate change will become increasingly costly to reverse. In part, this reflects the fact that damages are caused by the stock of GHGs (and not by the annual flow), but it also reflects the risks associated with irreversible thresholds and discontinuities. The Review therefore asks the question: what are the relatively certain and reversible costs that we must incur in the short term, in order to avoid potentially large, uncertain and irreversible damages in the future?

Some have argued that our estimates of mitigation cost are unrealistically low. This seems intuitively odd. For one thing, primary energy costs in most developed economies amount to around 3%–4% of GDP (a little higher in some industrialising countries). This means that anything less than a 25% increase in primary fuel costs is unlikely to raise costs above a maximum 1% of GDP. In practice, opportunities for energy efficiency and substitution to alternative processes and technologies are likely to keep the cost associated with such a rise in energy prices much lower than that (see Chapter 11). It was no surprise that our estimates of cost were reaffirmed by the IEA in their *World Energy Outlook 2006* (IEA, 2006), published after our Review. But it is worth re-emphasising that 1% of GDP is not a trivial amount. It represents a very significant change in the pattern of energy investment, in line with the replacement cycles of capital stocks, towards low-carbon energy technology.

Finally, it is important to understand precisely what 1% of GDP refers to. Many criticisms to the effect that costs could be higher in a world of poorly applied or poorly coordinated policy are not inconsistent with our assessment. This cost estimate applies in a fairly efficient world, where early action is taken. It reflects the likely costs under a flexible, global policy, employing a variety of economic instruments in cost-effective ways to control emissions of a broad range of GHGs. It would require clear, long-term price signals and policy frameworks that encourage technological

[17] The length of the irreversibility of mitigation costs would be dictated primarily by the lifetimes of capital equipment, infrastructure and networks. There is no reason to believe that inertia in endogenous factors such as knowledge, resulting, say, from a switch to alternative energies, would yield long-term net costs to society or inhibit long-term growth.

innovation. In the absence of these factors, or were action to be delayed or restricted to a limited number of countries, the costs would be significantly higher. By contrast, many of the higher cost estimates are generated through pessimistic assumptions about technological innovation and/or little substitution across production, inputs and technologies. This means that sectors continue to supply output at ever increasing costs as mitigation intensifies—an assumption that seems entirely out of line with human experience.

Studies that estimate costs under inappropriate policy in a particular country, or without the possibility of substitution across sectors, or which omit to model learning and innovation, will tend to estimate higher costs. This is so particularly later in the century, when it becomes necessary to tackle some of the more technologically complex sectors, such as transport, and to ensure very widespread deployment of zero-emissions power technologies, for which early investment in research, development and innovation is likely to be required.

More than $70bn (£36bn) of new money was invested globally in clean or renewable energy or clean technology last year. That constitutes a 43% increase on the year before, with considerable investments in solar, wind and biofuels. Goldman Sachs alone has invested $1.5bn in alternative energy in the past year. Financial-market estimates suggest that the global solar market will grow by 20%–30% every year in the next few years. This sector currently provides 4 gigawatts (GW) of energy worldwide, and it is expected that California alone will produce 3 GW of energy from solar in the next 10 years. By 2010, perhaps 10%–20% of this could be 'thin film', a technology that drastically cuts the amount of silicon needed (making it in some cases silicon-free) and so has the potential to substantially reduce cost. Companies investing in this include First Solar, Mitsubishi Heavy Industries, and Shell. There is now the very real prospect of rapid advances in solar technologies (using new materials), nano-battery storage and cellulosic biofuels.

Model estimates of global mitigation costs do not provide information on the incidence or distribution of these costs. Potential costs are likely to be higher in economies that are energy intensive, and lower in those that are less energy intensive or happen to be endowed with viable low-carbon alternatives. Since emissions are a global problem, there will need to be transfers to ensure that the lowest-cost mitigation strategies are pursued across the globe. In the second half of the Review (Parts IV and VI) we examine policy and in Part VI investigate possibilities for international carbon-trading markets to underpin these transfers

5. Putting it all together and the social cost of carbon

The Stern Review makes the case for strong and urgent action in three ways: a 'bottom-up' approach, comparing estimates of the disaggregated impacts of climate change with the costs of specific mitigation strategies; a 'model-based' approach, taking account of interactions in the climate system and the global economy; and a 'price-based' approach, comparing the marginal costs of emission abatement with the marginal damage cost or 'social' cost of emissions. Because CO_2 is the most significant human-generated GHG by overall volume and thus by overall warming, we often speak of the social cost of carbon.

The concept of the social cost of carbon is elaborated in Chapter 13 of the Review. Essentially, it is a measure of the impact of emitting an extra unit of carbon at any particular time on the present value (at that time) of expected well-being or utility, expressed in terms of a numéraire such as current consumption.[18] The concept is useful in three ways. First, it conveys the message that each and every GHG emission imposes a cost on society. Second, it can give some indication to policy-makers of the price that should be charged to those who emit GHGs. Third, by comparing it with the marginal cost of reducing emissions, it can help to demonstrate that there are net benefits of action to mitigate climate change.

However, the concept has its difficulties in definition, quantification and application, some of which we describe here. First, it is path dependent. That is to say, the marginal damage cost of carbon today depends on what happens to the stock of GHGs in the future, for as long as the gas emitted today stays in the atmosphere. To take a view on the social cost of carbon today means taking a view on the path of GHG emissions in the future, which depends on what policies are adopted and how the many uncertainties about climate-change impacts are resolved.

A second difficulty stems from the relationship between the social cost of carbon as defined here and the carbon prices (or tax rates) that would

[18] Thus a social cost of, say, £20 per tonne means that emitting an extra tonne of carbon today has the same impact on society's expected welfare as reducing consumption by £20 today.

be necessary to reduce emissions and lead to stabilisation. As the Review explained in Chapters 2 and 14, under some assumptions, setting a tax rate or price of a pollutant equal to its social cost ensures that the most desirable degree of abatement is obtained. The anti-pollution policy can be successfully 'decentralised'—brought about by the actions of many individuals without any further co-ordination.

But a range of strong assumptions are necessary to conclude that policymakers should put into effect a price of carbon equal to the social cost of carbon, because carbon pricing may induce other changes in behaviour, beyond emission reductions. Setting a carbon tax or price may lead to a fall in the prices of hydrocarbon energy sources, as the scarcity rents of exhaustible natural resources are reduced. This could help to undermine the effectiveness of climate-change policies (e.g. Sinclair, 1994). Similarly, setting a carbon tax or price may change the incentives to undertake research and deploy new technologies. The calculation of a social cost of carbon also entails choosing how to discount uncertain future climate-change impacts. If the policy-maker uses a lower discount rate than the households and firms making decentralised decisions, then the latter will not undertake as much investment in carbon abatement as the policy-maker would, were s/he responsible for the investment decision. The persistence of capital-market imperfections can complicate policy generally.

In the Review, we explained the importance of risk and uncertainty in setting pricing policy. Thus we argued that it was important, from the perspective of the economics of risk, to set quantitative goals for stabilisation and thus plan for (a corridor of) paths to be consistent with that stabilisation. Market prices for carbon would follow from these quantitative goals. The market mechanisms would help control costs of mitigation. Calculations of the social cost of carbon would then be used to assess whether market prices were developing in a reasonable way, and could contribute to revisions of quantitative targets. But, in this approach, social cost of carbon calculations would be checks on policy and not the prime drivers of policy. They are weak instruments in the latter role, since such calculations are so sensitive to assumptions.

It is plain, therefore, that there is no short cut to setting a carbon tax or price, which reinforces the case made in the Stern Review for a pragmatic approach. Decisions on the future trajectory of GHG emissions cannot be solely based on a formal 'optimisation' exercise, whereby the marginal abatement cost is set equal to the marginal damage cost, period by period. Critiques of the Review seem to overlook this point and may have laid far too much emphasis on particular calculations of the social cost of carbon.

6. Conclusion

Climate change presents a very complex challenge with a large number of dimensions, interactions, and dynamic feedbacks between separate impacts. It is important for policy-makers to find a structured way to think about the problem. No-one can predict specific impacts many years into the future with certainty, but scientists and economists now have a more comprehensive grasp of the risks and are providing policy-makers with informed judgements about the likely direction and scale of the changes. It is reasonable to ask policy-makers to take a strategic view of these threats, and to consider whether it is acceptable to incur certain costs now to reduce those risks. Of course, it is neither worth incurring any cost to reduce small risks, nor to make only a small reduction in the probability of very large risks. So policy-makers must make a judgement as to the likely costs and the cost-effectiveness of measures to reduce risk.

We looked at climate-change impacts in two ways. First, we carried out an analysis of its physical impacts around the world, most prominently on temperature, the water cycle, sea-level rise, and extreme weather events (in Chapters 1, 3, 4 and 5 of the Review). These are the core of our case for strong action relative to inaction. The case for action can be made independent of a precise estimation of the size of the impacts—provided the description of the impacts strongly suggests that they are larger than the costs of inaction. Because of the flow-stock process in accumulating GHGs in the atmosphere, delay can only add to the costs. Second, we conducted an analysis of the way in which formal economic modelling tries to aggregate and value those impacts (in Chapter 6). It is important to consider the evidence from the two approaches together. The formal economic modelling in Chapter 6 represents a highly aggregated attempt to capture these impacts and express them in terms of income/consumption. As such it is also highly simplistic and omits many of the crucial risks—it therefore has an inbuilt tendency to underestimate expected damages.

Much attention has nevertheless been focused on Chapter 6. Models require specific ethical frameworks and value judgements, and can be used to illustrate and explore the effect of different assumptions. But decision-makers require a richer understanding of the scale and nature of the risks involved in climate change. Different people will attach different weights to the different types of impacts. This is also a question of different ethical systems, not just of the narrow, usually utilitarian basis of cost–benefit economics. These should include rights to development of future generations, notions of sustainability and so on, which were explained in Chapter 2 of the Review, but not pursued at length. Attempts to model aggregate impacts should be accompanied by an analysis of specific regional and sectoral impacts (with probabilities attached). Better still is to start with detailed regional studies and to see formal aggregate analysis only as a supplement. That was the approach of the Review. It is the approach we strongly recommend for further research and policy analysis.

The question is often asked, why do our results place a higher valuation on climate-change impacts under BAU than many previous studies? The answer should by this stage be clear:

1. Our study takes on the latest science, including an explicit probabilistic assessment of high climate-change impacts;
2. We have included a broad range of impacts, from market impacts to environmental and health-related impacts, as well as catastrophic changes to the climate system, which scientists tell us are more probable at high temperatures. Many studies are limited to a subset of these impacts;
3. We have explicitly accounted for the economics of risk: risks and uncertainties are at the heart of concerns about climate change, but have hitherto often been finessed;
4. We have investigated ethical judgements about the valuation of future generations that are consistent and coherent in their treatment of individuals born at different times.

To omit any of these is to miss out an important element of the story. To include them will, however, provide a more realistic estimate of the cost of inaction. The question should rather be, 'why have previous studies produced such low impact estimates?' The answer corresponds to points 1 to 4 above: (i) they have mostly omitted to adequately employ the probabilistic results of recent science; (ii) they have tended to consider a narrow range of impacts, a product of focusing largely on 2°C–3°C warming, whereas we now know that there is a possibility of far higher temperatures; (iii) they have not used the economics of risk to the extent appropriate; (iv) they have not paid adequate attention to the underlying ethics. It is not surprising therefore that they have produced low estimates. There are some important exceptions to any one of these criticisms, but we think that the overall effect of insufficient attention on these fronts has given, on average, strong downward bias on damage estimates in the previous literature, taken as a whole.

Discounting has an important role to play in determining the value placed on climate-change impacts. This is a formal statement of the reality that those unconcerned about the welfare of future generations tend to be unconcerned about climate change. We argue that it is very hard to provide a plausible ethical justification for discounting human welfare on the basis of birth dates. But the sensitivity analysis makes clear that inter-generational welfare judgements are not the only—nor even the primary—factor driving the case for climate change; the treatment of risk and uncertainty and the extent to which the analysis and model embody progress in the scientific literature are of similar importance. The sensitivity analysis also illustrates that, extensive though it was, our aggregate analysis left out a number of risks and costs that would have increased our damage estimates still further. One reason for this is that science has recognised these risks, but not yet built them into its probability distributions (for example, various aspects of a changing carbon cycle via the oceans and forests).

The criticism that we chose assumptions in the formal modelling to give higher damages simply does not hold water. We left out intragenerational distribution from the formal modelling, and made no attempt to cost socially contingent impacts such as conflict and human migration, both of which have the potential to form a large part of the total estimated damages.

The question remains whether the results match our assessment of the likely disaggregated impacts under BAU, including our ability to cope with the outcomes, and our willingness to risk catastrophic events. Does a more than 5% loss in present global mean consumption accord with what we would expect from a disaggregated assessment of the list of likely social, economic and environmental damages associated with BAU? We feel that it does. The key judgement that has to be made, across all the analysis of whether strong and early action is justified, is whether the costs of such action are significantly smaller than the benefits of damages avoided. We think the answer is a clear 'yes'.

On the analysis of the cost of strong and early action, we concluded that mitigation costs under flexible, global policy could be around 1% of GDP (+/–3 percentage points), based on a wide range of modelling studies. The Review also considered expert opinion on technological possibilities and innovation, economic analysis on behavioural responses to relative price movements, and likely input and product substitution. Higher cost estimates are usually based on undue pessimism over policy or technological innovation, combined with historically unprecedented assumptions about a lack of substitution between products, inputs and technologies.

Over both the costs of inaction and the costs of action, our reflections on the analysis of the Review increase still further our confidence that our estimates are very much in the centre of any plausible range. Since publication of the Review, important international expert opinion has strongly endorsed our position. In particular, the IPCC's (2007) AR4 has reaffirmed our projections of climate change and assorted probabilities, while the IEA estimates are below our estimates, but near the centre of our range.

We argue that the costs of inaction far exceed the costs of action and that a sensible stabilisation range is in the region 450–550 ppm CO_2e. Anything higher than 550 ppm CO_2e is a very dangerous place to be, and the most severe damages can be avoided at affordable cost; anything lower than 450 ppm CO_2e is simply too expensive to realistically attain under current technologies and would not justify the damages avoided. Because the stock of GHGs is growing rapidly, the costs of meeting these targets rise for every year that action is delayed. This argues for urgent, strong and international action.

Those economists and others arguing for weaker policies, which point to concentrations of GHGs far higher than 550 ppm CO_2e, have a very difficult case to make. They must argue that a world with temperature rises above 5ºC relative to pre-industrial times is a world with which future generations can cope relatively easily. They have an intellectual obligation to convince us that such risks are very small. For their policies would appear to give clear risks of such an outcome. They should make such a case transparently and explicitly. Yet they rarely do.

Perhaps we should not be too surprised that so much energy has been expended in defending previous studies. But this is in our view mistaken and irresponsible. The positions of four or five years ago were taken before scientific evidence really gave us our first handle on probabilities, and they were derived using estimates of only a subset of all the climate impacts considered possible. They could not have articulated the case for action that is now emerging: their hands were tied. Probability analysis in both the science and the economics is in its early stages, but it surely points in the direction of action that we have described.

We have shown that climate change could have very serious impacts on growth and development. The costs of stabilising the climate are significant but manageable. There is still time to avoid the worst impacts of climate change, if we take strong action now. To delay now is to risk taking us to GHG concentrations that threaten irreversible changes to global geography, making readjustment in the future very costly.

Put simply, the economics of risk points to the following argument. If we take strong action now and spend around 1% of GDP annually on reducing GHG emissions, we will develop new technology, reduce other pollutants, promote energy efficiency, and reduce deforestation. These are likely to be valuable even given the remote chance that the science is wrong. If, on the other hand, we decide not to take strong action in the belief, for example, that the science might be wrong, and hope that we can adapt to whatever comes our way, we run a large risk arriving in a position, in which damage is severe and from which we cannot escape. The conclusion from the point of view of a sensible approach to risk is surely clear.

References

Ackerman, Frank, and Ian J. Finlayson. 2007. The economics of inaction on climate change: a sensitivity analysis. *Climate Policy* (forthcoming).

Anderson, Dennis. 2006. Costs and finance of carbon abatement in the energy sector. Paper prepared for the Stern Review.

Azar, Christian, and Kristian Lindgren. 2003. Catastrophic events and stochastic cost–benefit analysis of climate change. *Climatic Change* 56: 245–255.

Azar, Christian, and Thomas Sterner. 1996. Discounting and distributional considerations in the context of global warming. *Ecological Economics* 19 (2): 169–184.

Barker, Terry, Mahvash S. Qureshi, and Jonathan Koehler. 2006. *The Costs of Greenhouse-Gas Mitigation with Induced Technological Change: a Meta-Analysis of Estimates in the Literature*. Cambridge, UK: 4CMR, Cambridge Centre for Climate Change Mitigation Research, University of Cambridge.

Byatt, Ian, Ian Castles, Indur M. Goklany, David Henderson, Nigel Lawson, Ross McKitrick, Julian Morris, Alan Peacock, Colin Robinson, and Robert Skidelsky. 2006. The Stern Review: a dual critique. Part II—economic aspects. *World Economics* 7 (4): 199–229.

Carter, Robert M., C. R. de Freitas, Indur M. Goklany, David Holland, and Richard S. Lindzen. 2006. The Stern Review: a dual critique—Part I: the science. *World Economics* 7 (4):167–198.

Ceronsky, Megan, David Anthoff, Cameron Hepburn, and Richard S. J. Tol. 2005. Checking the price tag on catastrophe: the social cost of carbon under non-linear climate response. *Working Paper FNU-87*. Hamburg: Hamburg University and Centre for Marine and Atmospheric Science, Research Unit Sustainability and Global Change.

Dasgupta, Partha. 2006. Comments on the Stern Review's Economics of Climate Change. Available at www.econ.cam.ac.uk/faculty/dasgupta/STERN.pdf

Downing, Tom E., David Anthoff, B. Butterfield, Megan Ceronsky, Michael Grubb, Jiehan Guo, Cameron Hepburn, Chris Hope, Alistair Hunt, A Li, Anil Markandya, S. Moss, Anthony Nyong, Richard S. J. Tol, and Paul Watkiss. 2005. *Scoping Uncertainty in the Social Cost of Carbon*. London: DEFRA.

Fischer, C., and R. D. Morgenstern. 2005. Carbon abatement costs: why the wide range of estimates? *Discussion Paper 03–42*. Washington, DC: Resources for the Future.

Friedlingstein, P., P. Cox, R. Betts, L. Bopp, W. Bloh, V. Brovkin, S. Doney, M. Eby, I. Fung, B. Govindasamy, J. John, C. Jones, F. Joos, T. Kato, M. Kawamiya, W. Knorr, K. Lindsay, H. D. Matthews, T. Raddatz, P. Rayner, C. Reick, E. Roeckner, K-G Schnitzler, R. Schnur, K. Strassmann, S. Thompson, A. Weaver, C. Yoshikawa, and N. Zeng. 2006. Climate-carbon cycle feedback analysis: results from C4MIP model intercomparison. *Journal of Climate* 19: 3337–3353.

Gedney, N., P. Cox, and C. M. Huntingford. 2004. Climate feedback from wetland methane emissions. *Geophysical Research Letters* 31 (20): L20503.

Hepburn, Cameron J. 2006. Discounting Climate Change Damages: Working Note for the Stern Review.

HM Treasury. 2003. *The Green Book: Appraisal and Evaluation in Central Government: Treasury Guidance*. 3rd ed. London: Stationery Office.

Hope, Chris. 2006. The marginal impact of CO_2 from PAGE2002: an integrated assessment model incorporating the IPCC's five reasons for concern. *Integrated Assessment* 6 (1):19–56.

IEA. 2006. *World Energy Outlook 2006*. Paris: IEA.

IPCC. 2007. *Climate Change 2007: The Physical Science Basis: Summary for Policymakers*. Geneva: IPCC.

Jorgensen, Dale W., B. H. Goettle, and B. H. Hurd. 2004. *US Market Consequences of Global Climate Change*. Washington, DC: Pew Center on Global Climate Change.

Mastrandrea, Michael D., and Stephen H. Schneider. 2004. Probabilistic integrated assessment of "dangerous" climate change. *Science* 304: 571–575.

Meinshausen, Malte. 2006. What does a 2degC target mean for greenhouse gas concentrations? A brief analysis based on multi-gas emission pathways and several climate sensitivity uncertainty estimates. In *Avoiding Dangerous Climate Change*, edited by H.-J. Schellnhuber, W. Cramer, N. Nakicenovic, T. Wigley, and G. Yohe. Cambridge, UK: Cambridge University Press.

Mendelsohn, Robert O. 2006. A critique of the Stern Report. *Regulation* Winter: 42–46.

Mendelsohn, Robert O., W. N. Morrison, Michael E. Schlesinger, and Natalia G. Andronova. 1998. Country-specific market impacts of climate change. *Climatic Change* 45 (3–4): 553–569.

Murphy, J. M., D. M. H. Sexton, D. N. Barnett, G. S. Jones, M. J. Webb, and M. Collins. 2004. Quantification of modelling uncertainties in a large ensemble of climate change simulations. *Nature* 430 (7001): 768–772.

Nakicenovic, Nebojsa, and Rob Swart, eds. 2000. *Special Report on Emissions Scenarios: a Special Report of Working Group III of the Intergovernmental Panel on Climate Change*. Cambridge: Cambridge University Press.

Nordhaus, William D. 1994. *Managing the Global Commons: the Economics of Climate Change*. Cambridge, Mass: MIT Press.

———. 2006a. The economics of hurricanes in the United States. Paper read at Snowmass Workshop on Abrupt and Catastrophic Climate Change, at Snowmass, CO.

———. 2006b. A review of the Stern Review on the economics of global warming. Available at http://nordhaus.econ.yale.edu/SternReviewD2.pdf

Nordhaus, William D., and Joseph Boyer. 2000. *Warming the World: Economic Models of Global Warming*. Cambridge, Mass.: MIT Press.

Peck, S. C., and T. J. Teisberg. 1992. CETA: a model for carbon emissions trajectory assessment. *Energy Journal* 13 (1): 55–77.

Plambeck, Erica L., Chris Hope, and J. Anderson. 1997. The PAGE95 model: integrating the science and economics of global warming. *Energy Policy* 19: 77–101.

Quiggin, John. 2006. Stern and the critics on discounting. *Mimeo*: School of Economics and School of Political Sciences and International Studies, University of Queensland.

Rees, Martin J. 2003. *Our Final Century*. London: Heinemann.

Roughgarden, Tim, and Stephen H. Schneider. 1999. Climate change policy: quantifying uncertainties for damages and optimal carbon taxes. *Energy Policy* 27: 415–429.

Schellnhuber, Hans Joachim, Wolfgang Cramer, Nebojsa Nakicenovic, Tom Wigley, and Gary W. Yohe. 2006. *Avoiding Dangerous Climate Change*. Cambridge, UK: Cambridge University Press.

Sinclair, Peter J. N. 1994. On the optimum trend of fossil fuel taxation. *Oxford Economic Papers* 46: 869–877.

Smith, Joel B., Hans-Joachim Schellnhuber, and M. Monirul Qader Mirza. 2001. Vulnerability to climate change and reasons for concern: a synthesis. In *Climate Change 2001: Impacts, Adaptation, and Vulnerability: Contribution of Working Group II to the Third Assessment Report of the Intergovernmental Panel on Climate Change*, edited by J. J. McCarthy, O. F. Canziani, N. A. Leary, D. J. Dokken, and K. S. White. Cambridge, UK: Cambridge University Press.

Stern, Nicholas. 2006. *The Economics of Climate Change: The Stern Review*. Cambridge, UK: Cambridge University Press.

Tol, Richard S. J. 1994. The damage costs of climate change: a note on tangibles and intangibles, applied to DICE. *Energy Policy* 22 (5): 436–438.

———. 1999. Safe policies in an uncertain climate: an application of FUND. *Global Environmental Change* 9: 221–232.

———. 2002. Estimates of the damage cost of climate change—part II: dynamic estimates. *Environmental and Resource Economics* 21: 135–160.

———. 2003. Is the uncertainty about climate change too large for expected cost–benefit analysis? *Climatic Change* 56: 265–289.

Tol, Richard S. J., and Gary W. Yohe. 2006. A review of the Stern Review. *World Economics* 7 (4): 233–250.

Warren, R., C. Hope, M. Mastrandrea, R. Tol, N. Adger, and I. Lorenzoni. 2006. Spotlighting impacts functions in integrated assessment: research report prepared for the Stern Review on the Economics of Climate Change. *Working Paper 91*: Tyndall Centre for Climate Change Research.

APPENDIX
Addressing Confusions and Misconceptions about The Stern Review

The Stern Review is a large and detailed document that addresses a variety of complex scientific, ethical and economic issues. It is important that these are discussed and examined closely. It is unsurprising, therefore, that there has been some misunderstanding or misinterpretation of some aspects of the material. However, in some cases the errors have been careless, have not been followed up (and corrected) in consultation with us, and have been afforded early and misleading publicity. Here is a list of some of the most common misunderstandings relating to the economics.

Different business-as-usual (BAU) assumptions are made in assessing the impacts of climate change, compared with the costs of mitigation

The BAU assumptions used to assess the costs of mitigation are deliberately consistent with the BAU scenarios used to assess the impacts of climate change.[19] Our assessment of the impacts of climate change makes use of the IPCC *SRES* A2 scenario (Nakicenovic and Swart, 2000). In 2050, the A2 scenario predicts global annual CO_2 emissions of 60.4 Gt. This is very close to the BAU emission projections made in estimating the costs of mitigation. Annual BAU emissions in Chapter 7 reach 58 Gt CO_2 in 2050. The BAU emission projections underlying Anderson's (2006) resource-cost estimates in Chapter 9 lead to emissions of 60.9 Gt CO_2 in 2050.

The Stern Review has confused the units and quantities of CO_2 and CO_2e throughout the report

We have not confused CO_2 and CO_2e. The criticism we have received on this stems from a failure to realise that the pre-industrial level of CO_2e and CO_2 is actually the same.

The Stern Review uses different long-term growth assumptions in different chapters

It has been claimed that we are inconsistent in using two different long-term growth assumptions in different chapters (1.3% or 1.9%). In fact, the 1.3% annual growth figure quoted at the bottom of Box 6.3 is per capita, compared with the 1.9% aggregate growth figure referred to elsewhere in the Review. The 0.6% discrepancy is due to population growth.

[19] Except that, with respect to mitigation costs, we make projections up to 2075 only and, in some cases, only for fossil fuels—see Chapter 7, Footnote 8, Section 8.7 and also Figure 9.3 for details. We assume, essentially, that after 2075 costs remain constant around 1% of GDP.

Other IPCC *SRES* scenarios should have been applied instead

The IPCC states that all of its *SRES* scenarios are to be regarded as equally valid. This does not mean that they are all equally likely. Each is coherent and based on explicit and plausible (if not equally plausible) assumptions. But some scenarios invoke social and technological changes that are hard to describe as BAU. Our choice of the A2 scenario was based primarily on our own assessment of the most likely long-term developments in global GHG emissions. On the other hand, A2 has what would appear to be excessive population growth relative to productivity growth, so the sensitivity analysis we offer in this paper begins to explore alternative assumptions on population.

The science is out-of-date

Our summary of the science in Chapter 1 and our disaggregated analysis of impacts in Chapters 3–5 include references to many recent scientific papers. Climate predictions from the 2001 IPCC TAR are augmented, for example, by recent probabilistic evidence from the Hadley Centre (Murphy *et al.*, 2004) and from Meinshausen (2006), which offer an up-to-date grasp of probable warming. In Chapter 6, our high-climate scenario includes positive natural feedbacks drawn from modelling studies published within the last 3 years (e.g. Friedlingstein *et al.*, 2006; Gedney *et al.*, 2004). Over the course of the Review, we talked to many of the leading experts drafting chapters for the IPCC's 2007 AR4. Their insights influenced our assumptions and explain the similarities between our findings and the IPCC findings published subsequent to the Review.

The PAGE Integrated Assessment Model is biased—it is an outlier

See Section 3 of the main paper. On the contrary, PAGE2002 is designed to span the range of estimates put forward in previous studies, which is illustrated in Figure 2.

The Stern Review's formal economic modelling double-counts catastrophic risk (Tol and Yohe, 2006)

Tol and Yohe (2006) mistakenly presume that the risk of catastrophic climate change estimated by PAGE2002 is exactly calibrated on Nordhaus and Boyer's (2000) study, which is essentially expressed in terms of a certainty equivalent and thus already reflects some (though by no means the full) degree of uncertainty. But PAGE2002 is not calibrated on Nordhaus and Boyer (2000). Instead, it is calibrated on a wider range of evidence put forward in the IPCC TAR (see Hope, 2006, for a full account). The main paper above also puts forward the case that, due to a lack of data on underlying model uncertainty, PAGE2002 generally underestimates the uncertainty around the impacts of climate change. Thus the situation would more accurately be described as 'part-counting' risk: certainly not 'double-counting' it.

Adaptation was omitted from the assessment of climate-change impacts
One of the Review's six main parts (Chapters 18, 19 and 20) is devoted to adaptation. In our formal economic modelling, we also take explicit account of adaptation. Our central modelling scenario assumes that, in industrialised parts of the world, adaptation will reduce the impacts of climate change on market sectors of the economy by 90%, at all levels of warming. In lower-income regions—Africa, India and Southeast Asia, and Latin America—adaptation reduces market impacts by 50%, irrespective of warming. Worldwide, non-market impacts, primarily on human health and ecosystems, are reduced by 25% in value terms through adaptation, again at all levels of warming.

The Stern Review assumes a near-zero discount rate
This error stems from confusing the pure rate of time preference—which is nearly zero (0.1% per year)—with the overall social or consumption discount rate, which is the product of the elasticity of marginal utility of consumption—unity in our central case—and the growth rate of per-capita consumption.

The Stern Review won't say what discount rate is used
Because the discount rate depends on the growth rate of consumption per capita, and because the PAGE2002 model produces a range of different consumption growth paths, there is no unique discount rate. Indeed, it is a serious analytical mistake to think that there should be one discount rate.

It is impossible for the reader to understand calculations, while modelling methods are obfuscated
Chris Hope's PAGE2002 model is available and can be re-run on request. It is fully described in Hope (2006). Our social welfare valuation is fully specified and articulated in Box 6.3. Dennis Anderson's (2006) modelling work has been published on the internet and its technology resource data, assumptions and results are available on request too.

The Stern Review confuses income and consumption in its formal economic modelling
Saving is exogenous to the model of Chapter 6 such that—as pointed out in the text—consumption and income growth are identical. Where the effect of mitigation costs is analogous to an increase in a cost or price index (as would often be the case), the same percentage change applies to both consumption and income.

The modelling horizon for mitigation costs is incompatible with the modelling horizon for impacts, because the former is truncated in 2050
This is addressed in the main paper. Only the Anderson (2006) study was truncated in 2050, and we point out clearly that most studies suggest an average cost

of 1% into the medium term, but with growing uncertainty (see end of Section 10.3). Very few models of mitigation costs extend beyond 2100 because of uncertainties (see above). On the other hand, our comparison of costs of action and inaction assume that the 1% of GDP as mitigation cost carries on into perpetuity.

There is a mismatch in our approach to calculating the present value of mitigation costs compared with the present value of impact costs, because mitigation costs are not discounted, unlike impact costs

Our central estimate of an annual mitigation cost of 1% of GDP to follow a 550 ppm CO_2e stabilisation pathway is taken to be constant through time. It is not a present value and we do not calibrate impact costs in terms of present value. Indeed we have argued that we have to work in terms of the expected welfare integral, not marginal changes of the present-value kind. We compare a flow of mitigation costs (1% of GDP) with a percentage flow of climate-change impacts, calibrated via the balanced growth equivalent (BGE).

We accounted for risk aversion in estimating the costs of climate change, but not in estimating the costs of mitigation

The distribution of risks around the costs of mitigation is far narrower than that around the costs of climate change, so this would have been unlikely to change our results much in the near term.

Mitigation-cost estimates are drawn exclusively from (depending on the critique) either Anderson (2006) or Barker *et al.* (2006), both of which are biased downwards

Our central estimate of an annual cost of 1% of GDP through to the end of this century (with a range of +/–3% by 2050) is based on a very wide range of model estimates, including, for example, Stanford University's Energy Modelling Forum (EMF), those submitted to the US Climate Change Science Programme (CCSP) and the Innovation Modelling Comparison Project (IMCP). It does not stem from the use of one particular approach.

Mitigation cost estimates are implausibly low and out of line with other studies

Our estimate of likely mitigation costs was drawn from the literature and supported by our own analysis. Moreover, 1% of GDP is not a trivial amount and is in the middle of a range of the most respected modelling studies. It represents a very significant change in patterns of energy investment, in line with replacement cycles for capital stock, towards low-carbon energy technology. It is important to be clear what 1% of GDP reflects. It reflects likely costs under a flexible, global policy, employing a variety of economic instruments to control emissions of a broad range of greenhouse gases. It would require clear, long-term price signals and policy frameworks that encourage technological innovation. In the absence of these factors, or were action to be delayed, the costs could be significantly higher.

Transaction costs, capital-scrapping costs and planning costs have been understated or ignored in our estimate of mitigation costs

This criticism is made in relation to costs in the UK, up to 2020. However, our estimates are of global costs over this century. Our estimates also assume full policy flexibility with early action and long-term price signals. With action to utilise cheap mitigation options across the world, investment can take place in line with the capital-replacement cycle, keeping costs low. However, we do not dispute that inconsistent or delayed policy will lead to much higher costs for individual countries and/or the world as a whole.

The range of mitigation-cost estimates is so low that the implied benefit–cost ratio includes infinity, which is silly

Zero or negative costs of mitigation (which generate infinite benefit–cost ratios) have often been identified in the literature. Put simply, they allow for the entirely plausible possibility that emission abatement may yield efficiency gains and technological improvement, which can outweigh its costs: win–win policies. The apparent problem that Tol and Yohe (2006) identify arises only when one switches to expressing estimates as a benefit–cost ratio. That expression is unnecessary to the argument (and has limited analytical status). In any case, in choosing to compare our estimate of the cost of mitigation to that of EMF-21, they reveal in figure 1 of their critique that the EMF range encompasses zero cost as well.

Marginal abatements costs are falling over time according to Anderson (2006)

Marginal abatement costs are not falling through time in the Anderson study, they are rising. Average costs are falling. We pointed out in the Review (Box 9.6) that, although marginal costs are likely to rise through time, in line with a rising marginal damage cost of carbon, average costs may rise, fall or stay the same, depending on the rate of technological progress. This is fully in line with the academic literature on the evolution of costs.

Anderson (2006) biases mitigation costs downwards by ignoring capital costs

This is not true. The net costs of capital-stock turnover are included in the annualised costs.

Mitigation costs are understated, because rent is not counted as a resource cost

At a global level, rent is indeed merely a transfer, not a resource cost, as the transfers sum to zero. Monopoly rent produces allocative inefficiencies as a result of price distortions, but the rent itself is not a resource cost. We explicitly point out that the presence of such rents requires careful policy consideration (and model complexity) in addressing market distortions. We later stress the important role of carbon capture and storage in influencing rents associated with fuels of low marginal extraction cost.

Rebound effects (where energy demand rises as efficiencies increase) are ignored in Anderson (2006)

Rebound effects may be important and these are included in the broad modelling exercises that underpin our cost estimates. However, as we point out repeatedly, Anderson (2006) is not a behavioural model. This makes it a very simple and transparent means or baseline for cross-checking more complex model studies. To the extent that the general omission of behavioural responses biases the results, it will most likely do so in a positive direction—i.e. the cost of mitigation will be overestimated—because consumers and producers are assumed to be unable to substitute to cheaper technologies and processes.

Climate-change costs are not considered along stabilisation paths

In fact they are fully included in our estimate of the marginal damage cost of carbon in chapter 13. On a path to stabilise atmospheric GHG concentrations at 550 ppm CO_2e, we reported a preliminary estimate of the marginal damage cost (or ‘social’ cost) of carbon of around \$30/tCO_2. Stabilising at 450ppm CO_2e, our preliminary estimate was \$25/tCO_2e. These estimates indicate that some climate impacts are unavoidable, even on a stabilisation path. However, the impacts are much lower than on BAU, where our estimate of the social cost of carbon is in the region of \$85/tCO_2.

There was no optimisation analysis to determine stabilisation ranges

Such an analysis would demand too much of formal modelling and probabilistic forecasts, which project hundreds of years into the future. Indeed, modern public economics is replete with examples, where the characteristics of optimisation are profoundly determined by the particularities of unconsidered modelling assumptions (for example, the role of indirect and direct taxes is very sensitive to assumptions on separability in demand and supply functions, which are often untested). Nevertheless, in order to hone in on a stabilisation target range, a large section of Chapter 13 is dedicated to ‘backing-out’ implied marginal conditions (see Sections 13.5 and 13.6).