Centre for
Climate Change
Economics and Policy

An ESRC Research Centre

Endogenous growth, convexity of damages and climate risk: how Nordhaus' framework supports deep cuts in carbon emissions

Simon Dietz and Nicholas Stern

June 2014

Centre for Climate Change Economics and Policy
Working Paper No. 180

Grantham Research Institute on Climate Change and the Environment
Working Paper No. 159

Grantham Research Institute on Climate Change and the Environment

E·S·R·C ECONOMIC & SOCIAL RESEARCH COUNCIL

UNIVERSITY OF LEEDS

LSE THE LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

The Centre for Climate Change Economics and Policy (CCCEP) was established by the University of Leeds and the London School of Economics and Political Science in 2008 to advance public and private action on climate change through innovative, rigorous research. The Centre is funded by the UK Economic and Social Research Council and has five inter-linked research programmes:

1. Developing climate science and economics
2. Climate change governance for a new global deal
3. Adaptation to climate change and human development
4. Governments, markets and climate change mitigation
5. The Munich Re Programme - Evaluating the economics of climate risks and opportunities in the insurance sector

More information about the Centre for Climate Change Economics and Policy can be found at: http://www.cccep.ac.uk.

The Grantham Research Institute on Climate Change and the Environment was established by the London School of Economics and Political Science in 2008 to bring together international expertise on economics, finance, geography, the environment, international development and political economy to create a world-leading centre for policy-relevant research and training in climate change and the environment. The Institute is funded by the Grantham Foundation for the Protection of the Environment and the Global Green Growth Institute, and has five research programmes:

1. Global response strategies
2. Green growth
3. Practical aspects of climate policy
4. Adaptation and development
5. Resource security

More information about the Grantham Research Institute on Climate Change and the Environment can be found at: http://www.lse.ac.uk/grantham.

Endogenous growth, convexity of damages and climate risk: how Nordhaus' framework supports deep cuts in carbon emissions

Simon Dietz[1,2] and Nicholas Stern[1,3]

September 18, 2014

Forthcoming in the 125th anniversary issue of
The Economic Journal

Abstract

'To slow or not to slow' (Nordhaus, 1991) was the first economic appraisal of greenhouse gas emissions abatement and founded a large literature on a topic of great, worldwide importance. In this paper we offer our assessment of the original article and trace its legacy, in particular Nordhaus' later series of 'DICE' models. From this work many have drawn the conclusion that an efficient global emissions abatement policy comprises modest and modestly increasing controls. On the contrary, we use DICE itself to provide an initial illustration that, if the analysis is extended to take more strongly into account three essential elements of the climate problem – the endogeneity of growth, the convexity of damages, and climate risk – optimal policy comprises strong controls. To focus on these features and facilitate comparison with Nordhaus' work, all of the analysis is conducted with a high pure-time discount rate, notwithstanding its problematic ethical foundations.

Keywords: climate change, climate sensitivity, damage function, endogenous growth, integrated assessment

JEL codes: Q54

1 ESRC Centre for Climate Change Economics and Policy, and Grantham Research Institute on Climate Change and the Environment, London School of Economics and Political Science.
2 Department of Geography and Environment, London School of Economics and Political Science.
3 Department of Economics, London School of Economics and Political Science, and President of the British Academy.

We would like to thank Emanuele Campiglio, Antoine Dechezleprêtre, Baran Doda, David Greenaway, Tom McDermott, Elisabeth Moyer, Antony Millner, Bill Nordhaus and Bob Pindyck for helpful comments and discussions, and the editor, Rachel Griffiths. We would also like to acknowledge

the financial support of the Grantham Foundation for the Protection of the Environment and the Economic and Social Research Council. We alone are responsible for the content.

Email for correspondence: s.dietz@lse.ac.uk

1 To slow or not to slow

'To slow or not to slow' by Bill Nordhaus (1991) is a landmark in economic research. As the first analysis of the costs and benefits of policies to abate greenhouse gas emissions[1], it opened the profession to a new field of application – climate change. Its importance is partly illustrated by the number of times that it has been cited – on 1150 occasions according to Google Scholar; 398 times according to the narrower, journals-only measure in ISI Web of Knowledge.[2]

The context within which Nordhaus' paper was written helps us understand its contribution. While the basic science of the greenhouse effect was set out in the 19th century by Fourier, Tyndall and Arrhenius, discussions surrounding the possible role of humans in enhancing it – and therefore causing global warming and climate change – began in earnest in the 1970s. For at least a decade climate change remained largely a scientific/environmentalist's issue, debated in specialist conferences and networks (Agrawala, 1998). Indeed it is important to stress that the science of climate change was running years ahead of the economics (something that arguably remains the case today in understanding the impacts of climate change; Stern, 2013).

By the late 1980s, however, climate change was becoming both a policy issue and increasingly political. In 1988 the Intergovernmental Panel on Climate Change (IPCC) was established and in 1990 it published the first of its regular and influential *Assessment Reports* to member governments. In 1989 the first meeting of (22) Heads of State to discuss climate change was held in the Netherlands, and various other major international summits that year also put it on the agenda. Most OECD countries already had their first climate-change targets by 1990 (Gupta, 2010), for instance the European Community, as it was then, had pledged to stabilise its carbon dioxide emissions at 1990 levels by 2000. In 1992 virtually all countries signed up to the United Nations Framework Convention on Climate Change (UNFCCC) at a major summit on the environment and development in Rio de Janeiro, with its objective to achieve "stabilization of greenhouse gas concentrations in the atmosphere at a level that would prevent dangerous anthropogenic interference with the climate system" (Article 2).

Yet despite the obvious ecological risks of unmitigated climate change, the question remained whether the benefits of avoiding these risks would outweigh

[1] Shortly afterwards Bill Cline (1992) published what is generally considered to be the other foundational analysis of climate-mitigation benefits and costs.

[2] Both accessed on 24th March 2014. However, these citation counts likely understate the paper's legacy considerably, since many will instead cite later work that is based on it (see Section 2).

the perhaps substantial cost of cutting emissions.[3] This is the central question that 'To slow or not to slow' sought to tackle, by combining a simple model of social welfare and production with an externality from greenhouse gas emissions, in general equilibrium. This model took "existing models and simplified them into a few equations that are easily understood and manipulated" (p920), something that has become a hallmark of Nordhaus' work in the area. In summary the main components of the model are:

- A single equation of motion for the global mean temperature, which rises in response to the difference between the temperature that would be obtained in long-run equilibrium given the current atmospheric stock of greenhouse gases, and the current temperature;
- An equation of motion for the atmospheric stock of greenhouse gases, in which some fraction of current emissions adds to the stock, at the same time as some fraction of the current stock 'decays' by diffusing into the deep ocean;[4]
- A social welfare function that is the discounted sum over time of utility per capita;[5]
- Utility takes the form of the logarithm of consumption per capita of a single, aggregate good;
- Consumption per capita is given by (exogenous) output, less the total cost of abating emissions, and the total cost of climate change;
- A reduced-form abatement cost function, in which the total cost of abatement depends on global aggregate emissions and emissions abatement;
- Reduced-form damages, in which the total cost of climate change depends on global mean temperature, but where global mean temperature is an index of a wider set of climatic changes including changes in precipitation and sea level.

[3]There is a problem in using the language of benefit-cost analysis, if it is interpreted in its common and narrow, marginal, fairly undynamic way, and where risk is also treated narrowly. Climate-change policy raises major questions of the strategic management of potentially immense risks and where different paths will have different endogenous learning and discovery. This broader perspective is a major focus of this paper and should be central to economic research on the topic.

[4]To get an idea of the simplicity of the modelling framework, especially the science module, note that a fully-fledged atmosphere-ocean general circulation model such as that of the UK Hadley Centre would comprise hundreds of thousands of equations.

[5]There is little plausibility in moral philosophy for a social welfare function that is the sum across generations of the (discounted) utility per capita of each generation, irrespective of the number of people in a generation, unless population is constant. Adding the (undiscounted) total utility of each generation is essentially utilitarian. Pure-time discounting can be given a utilitarian interpretation if the discounting is based on the probability of existence as a function of time, and that becomes an exponential function in continuous time if the end of the world is the first event in a Poisson process.

This modelling framework has had a lasting influence on the field and indeed several elements of it still constitute the 'industry standard' today. The most notable example of this is perhaps the idea of reduced-form damages.

According to the model, Nordhaus found that an optimal cut in the current flow of global emissions of 11% relative to the base level should be made in a medium scenario (given a rate of pure-time preference of 1% per annum and 'medium' damages). In a 'high' scenario, with no pure-time discounting and 'high' damages, a cut of global emissions of around one third would be optimal. The concluding section of this paper lays out these results, without commenting on the plausibility of the various scenarios. Nonetheless, that the optimal emissions cuts were not more than one third implied that only modest targets could be supported by economic analysis of this kind, in comparison with some targets being discussed in the political arena. As the editor of the issue in which the paper appeared wrote, it is "certainly a sobering antidote to some of the more extravagant claims for the effects of global warming" (Greenaway, 1991, p903).

2 The DICE model framework

While it was very much the purpose of 'To slow or not to slow' to cast climate-change mitigation as a dynamic, investment problem, in which abatement costs could be paid up front, so that climate change could be avoided several decades into the future, the model dynamics were unsatisfactory – the economy was assumed to be in a so-called 'resource steady state', in which all physical flows are constant. Therefore we were asked to consider the setting as being the middle of the 21st century, when such conditions might plausibly hold (we can now see that this is highly unlikely). Optimal emissions abatement was calculated by evaluating a marginal change to the steady-state level (and thus the optimal cuts mentioned above were in the steady state). Time was still relevant though, because, while the change in abatement costs was instantaneous, the change in damage costs would be experienced only after a delay (Equations 7-9, p926).

Nordhaus himself was well aware of the shortcomings and indeed a preliminary version of a more fully dynamic model had already been presented at a workshop by the time 'To slow or not to slow' had been published. This new model was called DICE (for a 'dynamic integrated climate-economy' model) (Nordhaus, 1992, 1993b,a, 1994). Many elements of 'To slow or not to slow' could still be found in the original DICE model, including the equation of motion of the atmospheric stock of CO2, log utility, and reduced-form abatement and damage costs. But at the core was a Ramsey-Cass-Koopmans model of economic growth, allowing evaluation not only of the optimal steady state, but also of the optimal transition path. The social welfare function was modified to include population, with the objective becoming the (pure-time) discounted sum of total, instantaneous social utility, while a slightly more complex model of temperature change was also added. Once again, the results of the analysis with DICE pointed to modest emissions controls, modestly increasing over time

– from 10% initially to 15% in the later 21st century.

Since these first studies with the DICE model, it has become the pre-eminent integrated assessment model (IAM) in the economics of climate change. New versions have been published periodically (e.g. Nordhaus and Boyer, 2000; Nordhaus, 2008), and a regionally disaggregated model (RICE) was also developed (Nordhaus and Yang, 1996). However, to look only at Nordhaus' own studies with DICE is to hugely understate its contribution, because, by virtue of its simple and transparent unification of growth theory with climate science (not to mention Nordhaus' considerable efforts to make the model code publicly available), it has come to be very widely used by others. The uses to which it has been put are too numerous to cover in a comprehensive manner. Some of the more significant examples include: the introduction of induced innovation in the energy sector (Popp, 2004); explicit evaluation of optimal adaptation policy (de Bruin et al., 2009); consideration of uncertainty and learning (e.g. Kolstad, 1996; Keller et al., 2004); and treating consumption of material goods and environmental quality separately, thus allowing evaluation of relative price changes (Sterner and Persson, 2008).

Some of these extensions have challenged the broad conclusion that optimal emissions control is modest. And indeed it is important to stress two things. First, through his own updating of DICE, Nordhaus' position, as formalised in the model and its results, has shifted over the years towards stronger emissions reductions, albeit incrementally. Second, one can readily see in Nordhaus' writings an awareness of the limitations of IAMs like DICE. Nonetheless, it is fair to say the perception remains that an analysis of the costs and benefits of climate change in an IAM does not support strong emissions cuts, under standard assumptions. For instance, in the wake of the publication of the *Stern Review on the Economics of Climate Change* (Stern, 2007) (which in fact used an IAM other than DICE), it has been suggested that the difference in policy recommendations between the Review and other studies lies very largely in the specification of a low pure-time discount rate (Nordhaus, 2007), a rate that some have questioned.[6] A central purpose of the rest of this paper is to explore whether a recommendation of modest emissions reductions does indeed follow from using the DICE framework. We ask; can the framework support strong controls on emissions, if restrictive assumptions about growth, damages and climate risk are relaxed? These assumptions arguably lead to gross underestimation of the benefits of emissions reductions in DICE and other IAMs (Stern, 2013).

First, we incorporate endogenous drivers of growth and we allow climate change to damage these drivers. This is in stark contrast to the current generation of IAMs, which rests directly or indirectly on the Ramsey-Cass-Koopmans model, where the major source of growth per capita in the long run is exogenous improvements in productivity, but where climate change only impacts on

[6] A careful exploration of the strong basis in moral philosophy for low pure-time discounting is provided in Stern, (forthcoming a, b). In many IAM studies high pure-time discounting is introduced without much discussion.

current output.[7] There are compelling reasons for thinking that climate change could have long-lasting impacts on growth (Stern, 2013) and there is now an emerging body of empirical evidence pointing in this direction (e.g. Dell et al., 2012), even though climatic conditions in the recent past have been relatively stable compared with what we now have to contemplate.

Second, we assume that the damage function linking the increase in global mean temperature with the instantaneous reduction in output is highly convex at some temperature. Consideration of some of the science, for example on tipping points, leads us in this direction (see also Weitzman, 2012). By contrast, most existing IAM studies assume very modest curvature of the damage function. The DICE default is quadratic, and it is well known that with the standard values of the functions' coefficients an implausible 18degC or so of warming is required in order to reduce global output by 50 percent.[8]

Third, we allow for explicit and large climate risks. We do so by considering the possibility of high values of the climate-sensitivity parameter; that is, the increase in global mean temperature, in equilibrium, accompanying a doubling in the atmospheric concentration of carbon dioxide. We conduct sensitivity analysis on high values, but also specify a probability distribution reflecting the latest scientific knowledge on the climate sensitivity as set out in the recent IPCC report (IPCC, 2013). Its key characteristic is a fat tail of very high temperature outcomes that are assigned low probabilities. By contrast, most IAM studies have ignored this key aspect of climate risk by proceeding with a single, best guess value for the climate sensitivity, typically corresponding to the mode of the IPCC distribution. We note, linking the second and third points here, that the model temperature increase under business as usual a century or so from now of 3.5 or 4degC (see also IPCC, 2013) could be extremely damaging – this is not just a 'tail' issue.

Otherwise we remain faithful to the standard DICE framework, in order to make as clear as possible the difference that these three extensions make. Most notably, we retain its usual parameterisation of social values, where the rate of pure-time preference is 1.5% and the elasticity of marginal social utility of consumption is 1.5, so that with growth of consumption per capita of, say, 2%, the social discount rate would be 4.5%. We have written elsewhere about why we think it is inappropriate to posit such a high rate of pure-time preference (e.g. Stern, 2013, forthcominga,f) – and we return to explain why in Section 5 – but for the purpose of clarity of comparison we set aside our misgivings, concerning this and other features, in the modelling that comprises the core of this paper. More generally, there is a powerful case for arguing that this type of model, with one good and exogenous population, has very serious defects

[7] While a reduction in current output may impact future growth via reduced savings – for a given savings rate – we hypothesise that this effect is weak compared with direct reductions in the capital stock, and reductions in productivity. Fankhauser and Tol (2005) also find a weak impact of climate change on growth via savings, using DICE. They did not, however, consider that climate damages could work on the capital stock or on productivity.

[8] Nordhaus sees the specification of the damage function for warming above 3degC as a 'placeholder' (see Stern, 2013), but it is a placeholder that can have a powerful effect on the conclusions as we will see below.

in its ability to capture key aspects of a problem for which destruction of the environment and potential loss of life on a major scale are central.

3 Extending DICE

3.1 Endogenous growth

In standard DICE the production function is:

$$Y_t = F(K_t, L_t) = (1 - \hat{D}_t)(1 - \Lambda_t)A_t K_t^{\alpha} L_t^{1-\alpha} \tag{1}$$

where A_t is the exogenous element of total factor productivity (TFP) at time t, K is capital, L is labour and $\alpha \in (0, 1)$ is the capital exponent. $(1 - \hat{D}_t)$ is the standard DICE damage multiplier (see below for a definition) and the key point is that this is the only pathway through which climate change affects growth – directly by multiplication with gross output in each period. Λ represents emissions abatement costs. In all of our analysis we will maintain standard assumptions about Λ and L, detailed in Annex I alongside many other aspects of the model.

In (1) the long-run growth rate of output, ignoring for one moment the role of climate change, is given by the sum of the growth rates of A_t and L_t as in the standard Solow (1956) model. Climate-change damages $\hat{D}_t$ (and abatement costs Λ_t) affect the level of output in each period, which means that they can have two effects on the long-run growth rate of output. First, period-to-period changes in $\hat{D}_t$ can effectively change the long-run output growth rate. Second, depending on the rates of saving and capital depreciation, $\hat{D}_t$ can impact the long-run growth rate by affecting capital investment and in turn the stock of K in future periods.

Yet one of the central points of this paper is that this is a very narrow story of how climate change impacts on growth. We therefore consider two extensions to (1). Both are endogenous growth models, incorporating knowledge spillovers from the accumulation of capital by firms. And in both models, damages from a changing climate not only fall on gross output at a particular point in time, they also permanently reduce output possibilities at future points in time through their effect on endogenous determinants of growth.

A model of capital damages, and knowledge proportional to the capital stock

Our first growth model incorporates knowledge spillovers via the capital stock in the tradition of Arrow (1962), Romer (1986) and others. We combine this formulation with a partitioning of damages between output and capital. The production function becomes

$$Y_t = (1 - D_t^Y)(1 - \Lambda_t)A_t K_t^{\alpha+\beta} L_t^{1-\alpha} \tag{1.K}$$

where D^Y now denotes the damages that directly reduce annual output. In this model, we think of the economy as being composed of a number of firms, each making investments. Growth is driven in part by learning-by-doing, which in turn depends on each firm's net investment, so that when the firm's capital stock increases, so does economy-wide productivity. We also make the standard assumption in this tradition that knowledge is a pure public good. The elasticity of output with respect to knowledge is $\beta > 0$, so that the knowledge process has a productivity factor K^β. These assumptions have the effect of increasing the overall capital exponent to $\alpha + \beta$. We continue to assume an exogenous element of TFP A. This could be taken to represent elements of productivity not captured in knowledge spillovers, but we use it here principally for the narrower, instrumental purpose of calibrating (1.K) on (1) in the absence of climate-change damages and emissions abatement costs, thus achieving a controlled comparison of different production specifications.

We suppose there is further damage from climate change that reduces the capital stock, which we label D^K, so we obtain the following equation of motion of capital:

$$K_{t+1} = (1 - D_t^K)(1 - \delta^K)K_t + I_t \tag{2}$$

where $\delta^K \in [0, 1]$ is the depreciation rate on capital and $I_t = sY_t$ is investment, given savings rate s (see Annex I). In specifying D^K we have in mind the representation of two phenomena. First, D^K includes permanent, direct climate damage to the capital stock, for example if climate change increases the likelihood of storms and those storms damage infrastructure, or the abandonment of capital in coastal areas due to sea-level rise. Second, D^K could indirectly include broader impacts of climate change on productivity via the endogenous growth mechanism (1.K). One effect it could pick up is of a changing climate on the productivity of capital stocks, accumulated during a different and more stable climatic regime. For example, water supply infrastructure may become less productive given a long-run change in precipitation. Another could be that, if investment is increasingly diverted towards repair and replacement of capital damaged by extreme weather, it may produce fewer knowledge spillovers. Annex I contains further details of how, for our simulation work, we partition damages D between D^Y and D^K.

In sum, according to this model of growth and climate damages, some part of the instantaneous impacts of climate change falls on capital rather than output, so that this type of damage represents a permanent reduction in output possibilities in the future. Moreover since the economy's stock of knowledge is proportional to its stock of capital, the negative effect on future output possibilities is magnified.

A model of endogenous TFP and damages to TFP

One constraining feature of production functions like (1.K) is that, since knowledge is in one-for-one correspondence with the aggregate capital stock, it will

depreciate just as fast. If one considers a typical depreciation rate for economy-wide capital of 10% per year (indeed $\delta^K = 0.1$ in DICE), the implication is a rapid diminution of economy-wide knowledge over time. While the literature on measuring the returns to R&D investment points to annual depreciation of around 15% of private, firm-level R&D capital (see Hall et al., 2009, for a review), what we have here is a much broader construct of knowledge concerned with overall skills and know-how. Therefore we offer an alternative formulation of endogenous growth – new as far as we are aware – in which TFP is endogenous and depreciates more slowly than capital.

We revert to the standard production function, modelling TFP through a separate relation. The production function is hence:

$$Y_t = (1 - D_t^Y)(1 - \Lambda_t)\bar{A}_t K_t^\alpha L_t^{1-\alpha} \tag{1.TFP}$$

Capital and TFP have different dynamics. The equation of motion of the capital stock is simply given by

$$K_{t+1} = (1 - \delta^K)K_t + I_t \tag{2'}$$

Notice that in this specification we do not allow climate damages to impact the capital stock, although doing so would be straightforward by reverting to (2). The equation of motion of TFP is given by

$$\bar{A}_{t+1} = (1 - D_t^A)(1 - \delta_t^A)\bar{A}_t + a(I_t) \tag{3}$$

where δ^A is the *net* depreciation rate for productivity. We can think of δ^A as encapsulating both (a) depreciation of productivity through erosion or displacement of skills and know-how, and (b) implicit, autonomous growth of TFP, which captures among other things institutional innovations, beyond the scope of this model. Given these two effects, δ^A could in principle be negative, but here we shall assume it is positive and less than δ^K.

D^A is the part of damages that reduces productivity. It captures the productivity effects of climate change mentioned above. Annex I again explains how we partition D between D^Y and D^A.

$a(I_t)$ is a 'spillovers' function that converts the flow of capital investment in each period into a flow of knowledge externalities across activities as a whole. This means that the stock of TFP is augmented by knowledge spillovers, as well as changing over time according to the balance of depreciation and autonomous growth due to other factors, which is encoded in δ^A. In general assume $a' \geq 0$. More specifically, in order to again calibrate this model to standard DICE in the absence of climate damages and abatement costs, it is necessary to further assume $a' > 0$, $a'' < 0$, since in the standard DICE model the growth rate of TFP falls rapidly in the initial periods.[9] These properties can be satisfied by

$$a(I_t) = \gamma_1 I_t^{\gamma_2}$$

[9] Whether such concavity is theoretically or empirically plausible is not for this paper.

where $\gamma_1 > 0$ and $\gamma_2 \in (0, 1)$. Summing up, in this formulation some part of the instantaneous impacts of climate change falls on TFP, permanently reducing future output possibilities.

3.2 Convexity of damages

The standard DICE damage function is of a convenient reduced form that has come to be widely used in the field. Since we write the damage multiplier as $(1 - \hat{D}_t)$, write the standard damage function as

$$\hat{D}_t = 1 - 1/(1 + \pi_1 T_t + \pi_2 T_t^2) \tag{4}$$

where T is the global mean atmospheric temperature relative to the period just before the industrial revolution. The coefficients π_1 and π_2 are estimated by fitting the function on data points, which comprise the sum of underlying sectoral studies of climate damages at particular degrees of global warming (mostly $T = 2.5 - 3$degC), for example studies of crop losses and changing energy demand for space cooling and heating.[10] We should recognise, however, that these are 'quasi' data points, since $T = 3$degC has not been seen on the planet for around 3 million years and might lead to radical transformations in global climatic patterns. Making assumptions about the form of (4) is made still more difficult by the complete absence of evidence on aggregate impacts for $T \geq 3$degC. The quadratic form was originally selected largely for convenience[11], but it results in implausibly low damages at high temperatures (Stern, 2008; Weitzman, 2012). This has prompted Martin Weitzman (2012) to suggest the following modification:

$$D_t = 1 - 1/(1 + \pi_1 T_t + \pi_2 T_t^2 + \pi_3 T_t^{6.754}) \tag{4'}$$

where the coefficient π_3 and its corresponding exponent are together used to satisfy the assumption that, at $T = 6$, 50 percent of output is lost[12]. This is the functional form we use in this paper[13], but, in addition to Weitzman's calibration of π_3 we offer a second, alternative calibration such that $D_t = 0.5$ when $T = 4$. Science and impact studies tell us that, not only could we cross several

[10] Note that within this set of studies are some estimates of the money value of direct welfare losses due to climate change, for example impacts on health and the amenity value of the environment.

[11] Which is why Nordhaus himself describes such functions and the assumptions they embody about damages at different temperatures as "placeholders subject to further research" (Stern, 2013). However, we will see data points of 4, 5 or 6degC, if we are negligent and unlucky, within decades. Hence it makes sense to try different formulations as representing different possibilities, including of the extremely damaging circumstances the science suggests as possible.

[12] A quadratic function could not be made to simultaneously fit the existing data, while satisfying this additional assumption; it would give excessive damages for smaller temperature increases

[13] Elsewhere Dietz et al. (2007a,b,c); Stern (2007, 2008) we investigated models based on the PAGE IAM, in which damage was a power function of temperature. We examined the sensitivity of damages to the exponent of the power function up to a value of three.

key physical tipping points in the climate system by the time the 4degC mark is reached (Lenton et al., 2008), the impacts of such warming on the natural environment, economies and societies could be severe, with reason to believe in the risk of vast movements of population and associated conflict, unrest and loss of life (Stern, 2013). Global mean temperatures regularly exceeding 4degC above pre-industrial have likely not been seen for at least 10 million years (Zachos et al., 2008) and are within the range of difference between today and the peak of the last Ice Age, when large ice sheets covered northern Europe and North America (IPCC, 2013), radically influencing where people could be. Given the potential magnitude of transformation illustrated by this example, the assumption that $D_t = 0.5$ when $T = 4$ may be no less plausible, to put it cautiously, than assuming, as (4) does with the standard parameterisation, that $D_t = 0.04$ when $T = 4$, i.e. only 4% of output is lost as a result of temperatures not seen for 10 million plus years.

In our first growth model we partition damages as expressed in (4') between damages affecting output D^Y and those affecting capital D^K, while in our second model damages are partitioned between output and TFP as in (3). We do so in a similar way to Moyer et al. (forthcoming) and the procedure is described in detail in Annex I.

3.3 Climate risk

Our last extension to the basic framework involves the climate sensitivity parameter. We take two approaches here. First we explore high values of this parameter in sensitivity analysis. Second we replace its sure value with a probability density function (pdf). Climate sensitivity is a key factor in driving the change in temperature in DICE, as it is in many other simple climate models. Thus it is a natural example of large-scale risk. Others would be relevant too, such as the scale of damages for a given temperature increase, the scale of loss of life, and so on.

The equation of motion of temperature is given by:

$$T_t = T_{t-1} + \kappa_1 \left[F_t - \frac{F_{2\times CO_2}}{S} \left(T_{t-1} \right) - \kappa_2 \left(T_{t-1} - T_{t-1}^{LO} \right) \right] \tag{5}$$

where F_t is radiative forcing, $F_{2\times CO_2}$ is the radiative forcing resulting from a doubling in the atmospheric stock of carbon dioxide, S is the climate sensitivity, T^{LO} is the temperature of the lower oceans, κ_1 is a parameter determining speed of adjustment and κ_2 is the coefficient of heat loss from the atmosphere to the oceans. Calel et al. (forthcoming) contains a detailed explanation of the physics behind this equation.

In standard DICE $S = 3$degC. However, it has long been known that there is substantial uncertainty about S (Charney et al., 1979). Moreover investigations in recent years (as collected by Meinshausen et al., 2009) have tended to yield estimates of the pdf of S that have a large positive skew and in most

cases the right-hand tail can indeed be defined as 'fat'[14]. In the latest IPCC report (IPCC, 2013), a subjective pdf is offered that is the consensus of the panel's many experts. According to this distribution, S is 'likely' between 1.5 and 4.5degC, where likely corresponds to a subjective probability of anywhere between 0.66 and 1. It is 'extremely unlikely' to be less than 1C, where extremely unlikely indicates a probability of ≤ 0.05 , while it is 'very unlikely' to exceed 6degC, where this denotes a probability of ≤ 0.1. We thus choose values of $S \in \{1.5, 3, 6\}$ for sensitivity analysis.

For our stochastic modelling we fit a continuous pdf to these data, using the midpoints of the IPCC probability ranges. In doing so we face a choice over the type of function to fit. We performed a test of the fit of various functional forms, in terms of root-mean-square error, to the IPCC probability statements and found that the log-logistic function demonstrated the best fit among those we examined. The log-logistic function also has the advantage of having a tail of intermediate 'fatness' relative to other forms, thus, in this sense, it constitutes a middle-of-the-road assumption:

$$f(S) = \frac{a \cdot \left(\frac{S}{b}\right)^{a-1}}{b\left[1 + (\frac{S}{b})^{a}\right]} \tag{6}$$

where $a \approx 4.2$ and $b \approx 2.6$ are the shape and scale parameters respectively, giving mean S of 2.9, a standard deviation of 1.4 and the 95th percentile at 5.3.

It is worth emphasising, before moving on to the results, that there are other potentially significant sources of risk attending to the impacts of greenhouse gas emissions. Some of these are in the climate system – for instance the effective heat capacity of the oceans (Calel et al., forthcoming) – yet a focus on S captures the essence of physical climate risk in a clear and simple way. Other sources of risk relate to damages for any given temperature and could also be modelled with probabilities, were the evidence to justify doing so. However, as we have argued, the damage functional form and parameterisation are currently very poorly constrained by evidence and therefore it seems appropriate to instead pursue this, potentially very important source of risk, via a more simple sensitivity analysis on different functions as proposed in Section 3.2.

4 Results

4.1 Baseline

At the heart of this exercise is an investigation into the prospects for growth and damage in a changing climate. Figure 1 plots baseline consumption per capita – that is, in the absence of controls on carbon dioxide emissions imposed by a social planner – under various scenarios over the next two centuries. The upper panel plots the forecasts of the model with production (1.K) and damages from

[14]Where the density in the upper tail approaches zero more slowly than the exponential distribution.

climate change on the capital stock, while the lower panel plots the forecasts of the model with production (1.TFP), where TFP growth is endogenous and where climate change reduces TFP.

The 'standard' trajectory represents the forecast of the standard DICE model without the various extensions we are considering in this paper. The starting year is 2005. It is of course the same in both panels and notice immediately by how much consumption per capita increases in it, powered largely by exogenous productivity growth[15] – in 2205 it is more than 15 times the 2005 level. This is despite a large increase in the atmospheric stock of carbon dioxide and in the global mean temperature (discussed below). Without large assumed improvements in the exogenous element of TFP, the increase in per-capita consumption would be much smaller.

Changing the model of growth begins to yield more pessimistic forecasts, although it does not by itself qualitatively alter the tendency for the future to be much better off than the present. Under the model with capital damage, consumption/head in 2205 is 13.3 times higher than in 2005, while under the model of productivity damages it is 11.4 times higher. Since total damages D_t in (4) are the same in the two models, simply being partitioned differently between damages on output, capital and TFP (see Annex I), the larger effect in the model of productivity damages partly reflects the longer lasting impact of climate change in this model, where depreciation of productivity is slow compared with capital.

The divergence in forecasts is much more marked, however, when we layer on greater convexity of damages as in (4'). With Weitzman's (2012) calibration, consumption per capita grows much more slowly after 2150 in the model of capital damage, while in the model of TFP damage it peaks around 2150 before actually falling thereafter. By 2205 it is only 8.3 and 5.8 times higher respectively than today. If the damage function is set such that damages equivalent to 50% of global output are assumed to occur upon 4degC warming, the collapse in living standards is much stronger, with consumption/head peaking before the end of this century and ending up in both models around or below the present level in real terms. A similar forecast is generated by Weitzman damages, when we instead increase the climate sensitivity parameter S to 6degC, which has a probability, as described above, of up to 0.1 according to IPCC. The two growth models yield similar forecasts in these cases, demonstrating the diminished importance of growth assumptions when instantaneous damages are severe and or warming is very rapid.

Changes in the atmospheric stock of carbon dioxide and global mean temperature, which drive these growth prospects, are shown in Figures A.1 and A.2 respectively in Annex II. Baseline emissions will take the atmospheric stock of carbon dioxide to nearly 800ppm by the end of this century in all the scenarios considered. The stock continues to increase after 2100, but there is some feedback of climate damages on emissions, which works through the depressive

[15] With no growth in labour, the long-run output growth rate implied by (1) is simply that of exogenous TFP.

Figure 1: Baseline consumption per capita, 2005-2205. The upper panel corresponds to the model with capital damages and with knowledge proportional to the aggregate capital stock, while the lower panel corresponds to the model of endogenous TFP growth and TFP damages.

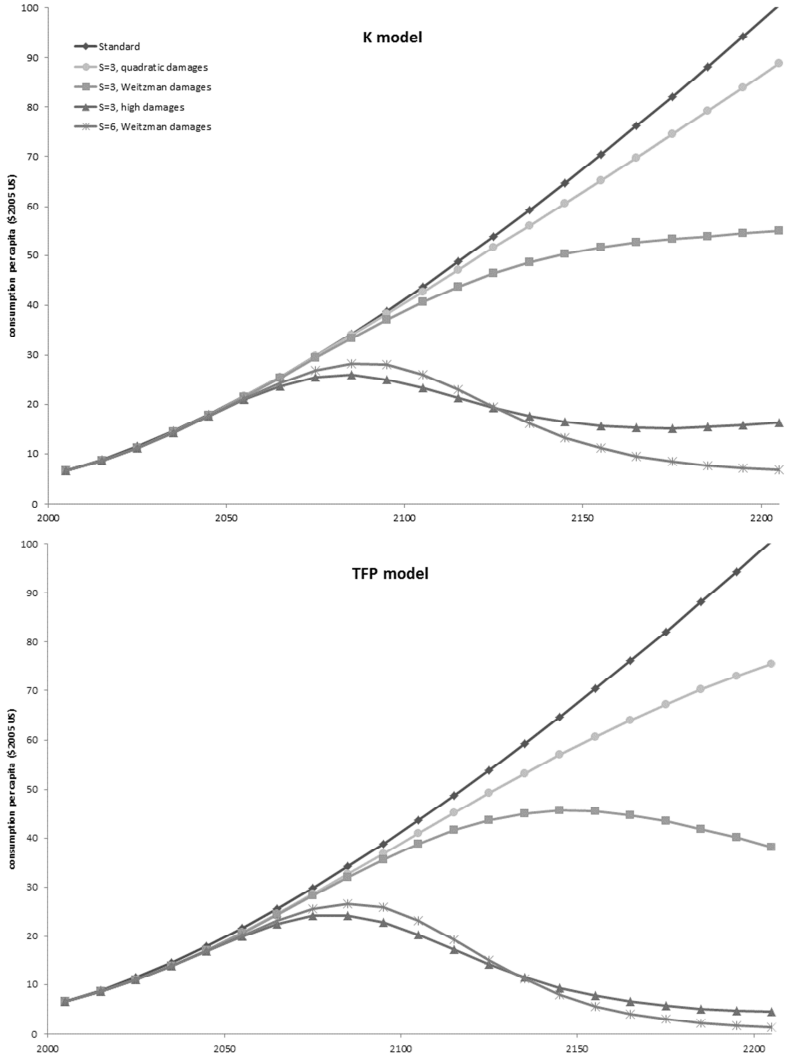

effect of climate damages on growth, and of growth on emissions. The principal determinant of global mean temperature is the value of the climate sensitivity parameter. With the typical central estimate of 3degC, the global mean temperature is forecast to be in the region of 3.5degC above the pre-industrial level by 2100, while if S=6 it could be more than 5degC above pre-industrial.

4.2 Optimal controls

We now move to examining the optimal controls on emissions, set by a social planner. As Annex I explains, the social planner's objective is to maximise the sum over time of discounted total utility by choosing a set of emissions control quantities and prices from 2015 until 2245, with a given abatement cost function (see Annex I). We present results covering the rest of this century. Table 1 lists the optimal emissions control rate (the percentage or fractional reduction in emissions from the baseline) under various scenarios, while Table 2 does the same for the optimal carbon price[16]. It is clear from the tables that modifying the growth model and the associated pathways through which climate change can affect the economy, as well as increasing the convexity of the damage function, and increasing the climate sensitivity, can significantly increase the optimal emissions control rate and the associated carbon price, both initially and throughout.

Let us focus on initial control quantities and prices – these give us something with which to compare current global policy efforts and debates. In standard DICE the emissions control rate, that is the percentage reduction in industrial carbon dioxide emissions, is 0.158 in 2015, with an associated carbon price of \$44/tC in 2005 prices (divide by roughly 3.7 to obtain estimates/tCO2, and multiply by c. 1.16 to bring up to 2012 prices[17]). If we switch from this standard model of exogenous growth to (1.K) with capital damages, the optimal emissions control rate rises to 0.213 (optimal carbon price = \$76/tC). Further extending this model to incorporate highly convex damages with Weitzman's (2012) parameterisation, it rises to 0.235 (optimal price = \$91/tC), while with our high damage function scenario it is 0.342 (optimal price = \$178/tC). When Weitzman damages are combined with a high climate sensitivity, the optimal control rate is 0.36, brought about by an optimal price levied at \$196/tC. Some caution should be exercised, however, in interpreting the relevance of these strong initial control rates and prices, because DICE, as a model of medium- and long-run dynamics, lacks adjustment costs, which could render such a rapid decarbonisation infeasible.

In the endogenous growth model (1.TFP) where instantaneous climate damages work on TFP as well as output, the increase in the controls is even stronger.

[16] Where the optimal carbon price is defined as the marginal cost of abatement at the optimal emissions level calculated. Whether it is reasonable to interpret this as a price depends on the convexity of the abatement cost curve, i.e. it depends on there being rising marginal costs. It has been contended that marginal costs do not rise, but these are issues for another paper.

[17] World Bank data on GDP deflator, from http://data.worldbank.org/indicator/NY.GDP.DEFL.KD.ZG, retreived on 22 November 2013

Table 1: Optimal emissions control rate, 2015-2105. $S = 3$ unless otherwise indicated; dge = damage.

	2015	2025	2035	2045	2055	2065	2075	2085	2095	2105
Standard	0.158	0.184	0.211	0.240	0.270	0.302	0.335	0.370	0.407	0.446
Capital	models									
Quadratic dge	0.213	0.289	0.356	0.424	0.495	0.565	0.636	0.706	0.777	0.848
Weitzman dge	0.235	0.322	0.401	0.484	0.568	0.650	0.730	0.805	0.875	0.944
Weitzman dge, S=6	0.360	0.494	0.619	0.751	0.883	1	1	1	1	1
High dge	0.342	0.471	0.591	0.709	0.814	0.901	0.970	1	1	1
Productivity	models									
Quadratic dge	0.272	0.365	0.444	0.528	0.614	0.702	0.792	0.882	0.974	1
Weitzman dge	0.290	0.392	0480	0.573	0.667	0.761	0.852	0.942	1	1
Weitzman dge, S=6	0.432	0.584	0.722	0.868	1	1	1	1	1	1
High dge	0.396	0.538	0.663	0.783	0.891	0.984	1	1	1	1

Table 2: Optimal carbon prices (2005 US$/tC), 2015-2105. $S = 3$ unless otherwise indicated.

	2015	2025	2035	2045	2055	2065	2075	2085	2095	2105
Standard	44.4	57.0	71.2	87.8	106.2	127.1	150.0	175.8	204.6	236.6
Capital	models									
Quadratic dge	76	129	182	245	316	393	476	563	656	752
Weitzman dge	91	156	226	310	405	506	609	711	812	912
Weitzman dge, S=6	196	337	495	684	895	1097	1074	1052	1032	1012
High dge	178	309	455	617	774	909	1017	1052	1032	1012
Productivity	models									
Quadratic dge	118	196	272	363	466	580	705	840	984	1012
Weitzman dge	133	222	313	420	541	670	806	945	1032	1012
Weitzman dge, S=6	271	456	653	888	1121	1097	1074	1052	1032	1012
High dge	233	393	559	738	911	1066	1074	1052	1032	1012

With quadratic damages, the optimal control rate on emissions is 0.272 with an associated carbon price of \$118/tC. Moving to Weitzman damages increases this to 0.29 (optimal carbon price = \$133/tC), while with our high damage function scenario the controls are respectively 0.396 and \$233/tC. When Weitzman damages are combined with a high climate sensitivity they are respectively 0.432 and \$271/tC. Notice for both growth models the marked rise in the carbon price when we move from Weitzman to high damages or from $S = 3$ to $S = 6$, which reflects convexity in the marginal abatement cost function. Nonetheless the same remarks regarding adjustment costs and their potential effect on the optimal controls apply here.

Figures A.3 and A.4 in Annex II show the consequences of the optimal controls for global mean temperature and the atmospheric stock of CO2. Compared with the baseline, it can be seen that maximisation of social welfare implies significant reductions in both climate variables. With Weitzman damages, the build-up of atmospheric CO2 is limited to 524ppm in the model of capital damages and 489ppm in the model of TFP damages, while with the more pessimistic parameterisation of the damage function the corresponding maximum concentrations are 459ppm and 444ppm. These numbers are broadly in line with the types of stabilisation concentrations recommended by many scientists. In sharp contrast, with standard DICE the optimal emissions controls allow atmospheric CO2 to rise throughout this century and peak at around 735ppm in the middle of the next century. The resulting warming depends on the climate sensitivity.

4.3 Optimal control under stochastic warming

Thus far we have computed the optimal controls, contingent on a set of point values of the climate sensitivity parameter, $S \in \{1.5, 3, 6\}$. A fuller specification of climate risk involves characterising a probability distribution over different values of S, as we described in Section 3, and solving for the optimal path of emissions controls. The planner's problem is specified as maximising expected social welfare, where expectations are formed before the first period commences and are not revised (see Annex I for further details of the optimal control problem).[18] Expectated values are formed in a Monte Carlo simulation, sampling (via the Latin Hypercube method) 500 times from $f(S)$ in (6).

Tables 3 and 4 report the optimal control quantities and prices respectively for the two growth models, each run with the various different damage functions. Since this exercise constitutes a fuller specification of climate risk, these might be considered our headline results. Notice that, comparing them with Tables 1 and 2, the effect of randomising S depends on the damage function – the optimal controls are higher under random S given Weitzman or high damages, but lower

[18] In line with much of the literature, we simplify the problem by omitting the possibility of learning about the climate sensitivity from observations obtained after the first period has commenced. So the planner must stick to optimal controls computed at the outset, a so-called open-loop control. Were it possible to learn about climate sensitivity from observations and to change policy settings in response – a closed-loop policy – the planner could of course achieve at least as high a level of social welfare, most probably much higher.

given quadratic damages. Remember that $f(S)$ in (6) is not a mean-preserving spread around $S = 3$. Rather, mean S is 2.9 and, as a distribution with a large positive skew, significantly more than half of the probability mass lies below the mean. When one bears in mind that what ultimately matters is the pdf of consumption per capita that results from $f(S)$, it should start to become clear that, when the damage function has modest curvature, the effect of randomising S on the optimal controls can be to lower them, but when the damage function has strong curvature the opposite is true, because the tail of high temperatures exerts an ever larger relative effect on consumption per capita, utility and social welfare.

Figures A.5 and A.6 in Annex II show the consequences of the optimal controls for the atmospheric stock of CO2 and global mean temperature respectively.[19] Figure A.5 shows that the optimal mean stock of atmospheric CO2 peaks in our endogenous growth models at no more than about 500ppm, and as little as 420ppm, depending on the growth model and damage function. These stock levels are well below those in the standard DICE model. Those combinations of growth model and damage function yielding higher climate impacts support a lower optimal stock. Compared with Figure A.3 we can see that the optimal stock is lower under random S than when $S = 3$. Figure A.6 shows that mean temperature is kept to a maximum of around 2degC except in two cases. First, in model (1.K) with capital damages, when the damage function is quadratic, mean warming peaks at around 2.5degC early next century. Second, in standard DICE mean warming peaks at c. 3.5degC. Notice the spreads around mean warming and in particular the very large 90% confidence interval around warming in standard DICE, where the 90th percentile reaches as much as 5.6degC. Optimal emissions controls in our extended models of DICE cut this tail of high temperatures significantly, due to their damaging consequences in the short and long run.

5 Conclusions

'To slow or not to slow' (Nordhaus, 1991) and its subsequent development into the dynamic DICE model have given us what seems to be a coherent and powerful framework for assessing the costs and benefits of climate-change mitigation. But it has in-built assumptions on growth, damages and risk, which together result in gross underassessment of the overall scale of the risks from unmanaged climate change (Stern, 2013). This criticism applies with just as much force to most of the other IAMs that DICE has inspired. The purpose of this article has been to show how these unrealistic assumptions might be relaxed

[19] Since the climate sensitivity is uncertain, so, obviously, is the change in the global mean temperature, and since this goes on to affect emissions via damages, there is also some uncertainty in the longer run about the atmospheric stock of CO2. Therefore both figures report mean values from the Monte Carlo simulation. In the case of the atmospheric stock of CO2 the uncertainty is very small (no more than 1ppm), but in the case of global mean temperature it is considerably larger. Therefore in the latter case we also show the 90% confidence interval, in 2205, from the Monte Carlo simulation to the right of the main chart.

Table 3: Optimal emissions control rate under random S, 2015-2105.

	2015	2025	2035	2045	2055	2065	2075	2085	2095	2105
Standard	0.158	0.184	0.211	0.240	0.270	0.302	0.335	0.370	0.407	0.446
Capital	models									
Quadratic dge	0.204	0.277	0.341	0.406	0.474	0.543	0.611	0.681	0.750	0.819
Weitzman dge	0.250	0.344	0.434	0.532	0.631	0.732	0.831	0.925	0.993	0.993
High dge	0.393	0.542	0.688	0.841	0.986	0.998	0.998	0.998	0.999	0.999
Productivity	models									
Quadratic dge	0.261	0.350	0.427	0.508	0.591	0.677	0.765	0.854	0.944	1
Weitzman dge	0.307	0.417	0.518	0.627	0.743	0.862	0.982	0.999	1	1
High dge	0.481	0.656	0.825	1	1	1	1	1	1	1

Table 4: Optimal carbon prices (2005 US$/tC) under random S, 2015-2105.

	2015	2025	2035	2045	2055	2065	2075	2085	2095	2105
Capital	models									
Quadratic dge	70	119	169	226	293	365	442	528	614	707
Weitzman dge	101	176	261	368	490	625	769	914	1019	999
High dge	229	399	598	838	1093	1092	1071	1050	1030	1010
Productivity	models									
Quadratic dge	110	181	253	338	435	543	664	792	930	1012
Weitzman dge	147	248	359	494	657	839	1040	1050	1032	1012
High dge	329	563	830	1146	1121	1097	1074	1053	1032	1012

and what would be the consequences of doing so, in terms of optimal emissions reductions and carbon prices, atmospheric concentrations of carbon dioxide and global mean temperature.

The first assumption we have relaxed is that the underlying drivers of economic growth are exogenous and unaffected by climate change. Instead we look at two models of endogenous growth, in which the damages from climate change affect the drivers of long-run growth, not just current output. The second assumption we have relaxed is that the damage function relating instantaneous climate damage to the increase in global mean temperature is only weakly convex. Instead we allow for the possibility that instantaneous damages increase rapidly, particularly once the global mean temperature reaches 4-6degC above the pre-industrial level. We suggest this representation is more plausible given the scale of change that such warming could bring; at the very least, simulations based on weak convexity should not dominate our attention as they have come to do. The third assumption we have relaxed is that the climatic response to greenhouse gas emissions is moderate and moreover is precisely understood. Very few, if any, commentators would explicitly claim that climate sensitivity is precisely understood, of course. Nonetheless most economic modelling is undertaken using only a single, central estimate of the climate sensitivity parameter, fixed in the centre of the distribution of available estimates from the science. We explore risk in this crucial parameter.

Overall, the scale of the risks from unmanaged climate change in this modelling framework is the convolution of these three extensions. We show that, with the models extended in this way, business-as-usual trajectories of greenhouse gas emissions give rise to potentially large impacts on growth and prosperity in the future, especially after 2100. Indeed these impacts are large enough to feed back into future emissions via reduced activity, but the feedback is too small and too late for the system to self-regulate. Thus optimal emissions control is strong and strongly increasing. As a guide, we find that these models suggest the carbon price in a setting of globally coordinated policy, such as a cap-and-trade regime or a system of harmonised domestic carbon taxes, should be in the range \$32-103/tCO2 (2012 prices) in 2015. It must be remembered that the DICE model lacks adjustment costs, so the high end of the range should be interpreted cautiously. On the other hand and potentially of great importance, we have, notwithstanding our extensions, omitted important risks in relation to the distribution of damages, which could give higher carbon prices. Within two decades the carbon price should rise in real terms to \$82-260/tCO2. Doing so would, according to the model, keep the expected atmospheric stock of carbon dioxide to a maximum of c. 425-500ppm and the expected increase in global mean temperature to c. 1.5-2degC above pre-industrial.

The paper is only a preliminary investigation, whose purpose was to illustrate or sketch the consequences of relaxing assumptions that have limited plausibility and possible large effects on policy conclusions. We have, for instance, restricted our attention to knowledge, accumulated through learning-by-doing, as the driver of long-run growth, though other sources of growth are important and other models might be deployed. Our exploration of the implications of

risk has, for the sake of clarity, been limited to the climate sensitivity, though other sources of great risk exist in the physical climate system, not to mention in the economy. The models that we do use require the choice of parameter values, about some of which there is currently very little relevant empirical evidence. Given slow rates of learning about some IAM parameters, this should be regarded as an endemic problem, however (Pindyck, 2013). It is not the case that the standard model parameters are well constrained, whereas the new parameters we introduce are not. Future work building on our framework should also pay attention to the costs of rapid adjustment to a low-carbon economy and possible limits to the speed of decarbonisation. This work will need to go well beyond the choice of parameter values to consider new model structures.

This has not been a paper about the sensitivity of results to pure-time discounting, or other parameters and structures relevant to discounting. As we found in the technical annex to Stern (2007) and in Dietz et al. (2007c; 2007a), lower pure-time discounting does indeed favour stronger and earlier action to curb emissions. Those results were from the 'PAGE' IAM (Hope, 2006), but we know from other work that this is also true of DICE (Nordhaus, 2007). We have argued elsewhere that careful scrutiny of the ethical issues around pure-time discounting points to lower values than are commonly assumed (usually with little serious discussion). Pure-time discounting is essentially discrimination by date of birth in the sense that a life, which is identical in all respects (including time patterns of consumption) but happens to start later, has a lower value. If, for example, the pure-time discount rate were 2%, a life starting 35 years later, but otherwise the same, would have half the value of a life starting now. The time horizon essential to a discussion of climate change makes careful examination of these ethical issues unavoidable. Preliminary calculations indicate that low pure-time discounting will significantly increase the optimal controls in this paper as well.

One cannot and should not expect a single model to capture all relevant issues and neither should we be able to resolve all difficulties within a single framework.[20] It is enough for a model to help raise and understand key aspects of a problem. This means that we should be grateful to Bill Nordhaus for providing one helpful vehicle. As it is expanded and different perspectives are brought in, including the possibility of major loss of life from climate change, then we would suggest the arguments for strong action will look still stronger.

References

Agrawala, S., 1998. Context and early origins of the intergovernmental panel on climate change. Climatic Change 39 (4), 605–620.

Arrow, K. J., 1962. The economic implications of learning by doing. The Review of Economic Studies 29 (3), 155–173.

[20] For example, the paradoxes of social choice theory can be better understood by broadening the philosophical perspective and are not easily resolved within the standard framework (Sen, 2009).

Calel, R., Stainforth, D. A., Dietz, S., forthcoming. Tall tales and fat tails: the science and economics of extreme warming. Climatic Change.

Charney, J. G., Arakawa, A., Baker, D. J., Bolin, B., Dickinson, R. E., Goody, R. M., Leith, C. E., Stommel, H. M., Wunsch, C. I., 1979. Carbon dioxide and climate: a scientific assessment.

Cline, W. R., 1992. The Economics of Global Warming. Peterson Institute.

de Bruin, K. C., Dellink, R. B., Tol, R. S., 2009. Ad-dice: an implementation of adaptation in the dice model. Climatic Change 95 (1-2), 63–81.

Dell, M., Jones, B. F., Olken, B. A., 2012. Temperature shocks and economic growth: Evidence from the last half century. American Economic Journal: Macroeconomics 4 (3), 66–95.

Dietz, S., Anderson, D., Stern, N., Taylor, C., Zenghelis, D., 2007a. Right for the right reasons: a final rejoinder on the stern review. World Economics 8 (2), 229–258.

Dietz, S., Hope, C., Patmore, N., 2007b. Some economics of dangerous climate change: Reflections on the stern review. Global Environmental Change 17 (3), 311–325.

Dietz, S., Hope, C., Stern, N., Zenghelis, D., 2007c. Reflections on the stern review (1): a robust case for strong action to reduce the risks of climate change. World Economics 8 (1), 121–168.

Fankhauser, S., Tol, R. S., 2005. On climate change and economic growth. Resource and Energy Economics 27 (1), 1–17.

Greenaway, D., 1991. Economic aspects of global warming: Editorial note. The Economic Journal 101 (407), 902–903.

Gupta, J., 2010. A history of international climate change policy. Wiley Interdisciplinary Reviews: Climate Change 1 (5), 636–653.

Hall, B. H., Mairesse, J., Mohnen, P., 2009. Measuring the returns to r&d. Tech. rep., National Bureau of Economic Research.

Hope, C., 2006. The marginal impact of co2 from page2002: An integrated assessment model incorporating the ipcc's five reasons for concern. Integrated Assessment 6 (1).

IPCC, 2013. Working Group I Contribution to the IPCC Fifth Assessment Report: Summary for Policymakers. In: Climate Change 2013: The Physical Science Basis. IPCC.

Keller, K., Bolker, B. M., Bradford, D. F., 2004. Uncertain climate thresholds and optimal economic growth. Journal of Environmental Economics and Management 48 (1), 723–741.

Kolstad, C. D., 1996. Learning and stock effects in environmental regulation: the case of greenhouse gas emissions. Journal of Environmental Economics and Management 31 (1), 1–18.

Lenton, T. M., Held, H., Kriegler, E., Hall, J. W., Lucht, W., Rahmstorf, S., Schellnhuber, H. J., 2008. Tipping elements in the earth's climate system. Proceedings of the National Academy of Sciences 105 (6), 1786–1793.

Mankiw, N. G., Phelps, E. S., Romer, P. M., 1995. The growth of nations. Brookings papers on economic activity 1995 (1), 275–326.

Meinshausen, M., Meinshausen, N., Hare, W., Raper, S. C., Frieler, K., Knutti, R., Frame, D. J., Allen, M. R., 2009. Greenhouse-gas emission targets for limiting global warming to 2c. Nature 458 (7242), 1158–1162.

Mirrlees, J. A., Stern, N. H., 1972. Fairly good plans. Journal of Economic Theory 4 (2), 268–288.

Moyer, E., Woolley, M., Glotter, M., Weisbach, D. A., forthcoming. Climate impacts on economic growth as drivers of uncertainty in the social cost of carbon. Climatic Change.

Nordhaus, W. D., 1991. To slow or not to slow: the economics of the greenhouse effect. Economic Journal 101 (407), 920–937.

Nordhaus, W. D., 1992. An optimal transition path for controlling greenhouse gases. Science 258, 1315–1315.

Nordhaus, W. D., 1993a. Optimal greenhouse-gas reductions and tax policy in the "dice" model. The American Economic Review 83 (2), 313–317.

Nordhaus, W. D., 1993b. Rolling the "dice": an optimal transition path for controlling greenhouse gases. Resource and Energy Economics 15 (1), 27–50.

Nordhaus, W. D., 1994. Managing the Global Commons: the Economics of Climate Change. MIT Press, Cambridge, MA.

Nordhaus, W. D., 2007. A review of the "stern review on the economics of climate change". Journal of Economic Literature, 686–702.

Nordhaus, W. D., 2008. A Question of Balance: Weighing the Options on Global Warming Policies. Yale University Press.

Nordhaus, W. D., Boyer, J., 2000. Warming the World: Economic Models of Global Warming. MIT Press (MA).

Nordhaus, W. D., Yang, Z., 1996. A regional dynamic general-equilibrium model of alternative climate-change strategies. The American Economic Review, 741–765.

Pindyck, R. S., 2013. Climate change policy: What do the models tell us? Journal of Economic Literature 51 (3), 860–872.

Popp, D., 2004. Entice: endogenous technological change in the dice model of global warming. Journal of Environmental Economics and Management 48 (1), 742–768.

Romer, P. M., 1986. Increasing returns and long-run growth. The Journal of Political Economy, 1002–1037.

Sen, A. K., 2009. The Idea of Justice. Harvard University Press.

Solow, R. M., 1956. A contribution to the theory of economic growth. The Quarterly Journal of Economics 70 (1), 65–94.

Stern, N., 2007. The Economics of Climate Change: the Stern Review. Cambridge University Press.

Stern, N., 2008. The economics of climate change. American Economic Review: Papers and Proceedings 98 (2), 1–37.

Stern, N., 2013. The structure of economic modeling of the potential impacts of climate change: grafting gross underestimation of risk onto already narrow science models. Journal of Economic Literature 51 (3), 838–859.

Stern, N., forthcominga. Ethics, equity and the economics of climate change. paper 1: science and philosophy. Economics and Philosophy.

Stern, N., forthcomingb. Ethics, equity and the economics of climate change. paper 2: economics and politics. Economics and Philosophy.

Sterner, T., Persson, U. M., 2008. An even sterner review: Introducing relative prices into the discounting debate. Review of Environmental Economics and Policy 2 (1), 61–76.

Weitzman, M., 2012. Ghg targets as insurance against catastrophic climate damages. Journal of Public Economic Theory 14 (2), 221–244.

Zachos, J. C., Dickens, G. R., Zeebe, R. E., 2008. An early cenozoic perspective on greenhouse warming and carbon-cycle dynamics. Nature 451 (7176), 279–283.

Annex I: extended model description

Here we offer an extended description of the DICE model, focusing on the major model equations and in particular on our modifications. Even more detail can be found on Nordhaus' model website at http://www.econ.yale.edu/~nordhaus/homepage/index.html. Our analysis is based on the 2010 version of the model.

The model represents a social planner maximising a classical utilitarian objective functional by choosing the rate of control of industrial carbon dioxide emissions:

$$\max_{\{\mu_t\}_{t=1}^{Tmax}} W = \sum_{t=0}^{Tmax} u(c_t) L_t (1+\rho)^{-t}$$

where $\mu \in [0,1]$ is the emissions control rate, $u(c_t)$ is the instantaneous social utility of consumption per capita at time t and $\rho = 0.015$ is the utility discount rate. Note that c is not only time-dependent as the above equation implies, it is also state-dependent when we undertake stochastic modelling. We suppress notation of state-dependence for simplicity; bear in mind that, when running the model with a random parameter (the climate sensitivity S in equation 5), we take the expectation of social welfare. $Tmax$ is the terminal period, which is 2595. The model proceeds in time steps of ten years from 2005, so appropriate interpretations must be made in considering the various equations of motion. Notice that, since 2005 is in the past, our first control period is $t = 1$, i.e. 2015.[21]

The utility function is iso-elastic,

$$u(c_t) = \frac{c_t^{1-\eta}}{1-\eta}$$

where η is the elasticity of marginal social utility of consumption and is set to 1.5 to allow comparison with standard DICE.

As set out in the main body of the paper, we explore two alternative production functions:

$$Y_t = (1 - D_t^Y)(1 - \Lambda_t) A_t K_t^{\alpha+\beta} L_t^{1-\alpha} \tag{1.K}$$

$$Y_t = (1 - D_t^Y)(1 - \Lambda_t) \bar{A}_t K_t^{\alpha} L_t^{1-\alpha} \tag{1.TFP}$$

The capital elasticity $\alpha = 0.3$, while the elasticity of output with respect to knowledge, $\beta = 0.3$ (Mankiw et al., 1995).

The equation of motion of capital in model (1.K) is

$$K_{t+1} = (1 - D_t^K)(1 - \delta^K) K_t + I_t$$

where $\delta^K = 0.1$, while in model (1.TFP) we simply drop $(1 - D^K)$.

[21] In fact we need only solve μ_t from 2015 to 2245 inclusive, since DICE assumes that from 2255 onwards $\mu_t = 1$, because a zero-emissions backstop energy technology becomes competitive.

Output is either consumed or invested,

$$Y_t = C_t + I_t$$

where $C_t = c_t L_t$ is aggregate consumption and $I_t = sY_t$, where $s = 0.23$ is the savings rate (calibrated to long-run average optimal savings in standard DICE, absent climate damages and emissions abatement costs and, in principle, some private inter-temporal objectives). We specify exogenous, constant savings in order to capture in a simple way the second-best context implied by fitting our models of endogenous growth to current macroeconomic data. In growth models with knowledge spillovers, the savings rate chosen by a planner will be greater than the savings rate emerging from a decentralised equilibrium of firms and households, because the marginal private return to investment does not include the spillovers.

A more elaborate analysis would permit households to choose their optimal savings rate in equilibrium with firms' private marginal product of capital (in response to the planner's emissions controls), but it is worth noting that, in standard DICE, endogenising the savings rate has been shown to make little difference to the optimal policy (Fankhauser and Tol, 2005, and Nordhaus' lab notes on DICE), so our simplification is unlikely to matter.[22] In any case, whether households are currently taking into account the effects of climate policy on future consumption prospects when choosing how much to save is unclear.

In model (1.TFP) productivity is endogenous and its equation of motion is

$$\bar{A}_{t+1} = (1 - D_t^A)(1 - \delta^A)\bar{A}_t + \gamma_1 I_t^{\gamma_2} \tag{3}$$

where $\delta^A = 0.01$ is the rate of depreciation of the stock of TFP, while $\gamma_1 \approx 0.0003$ and $\gamma_2 \approx 0.373$ are parameters of the spillovers function. γ_1 and γ_2 are calibrated so that output in (1.TFP), in the absence of climate damages and emissions abatement costs, is the same as in standard DICE. In model (1.K) TFP is an exogenous time series, so (3) does not apply.

The climate damage function is

$$D_t = 1 - 1/(1 + \pi_1 T_t + \pi_2 T_t^2 + \pi_3 T_t^{6.754}) \tag{4'}$$

where $\pi_1 = 0$ and $\pi_2 \approx 0.00284$ throughout. $\pi_3 = 0$ when we compute results for the standard setting (i.e. $D_t = \hat{D}_t$), $\approx 5.07 * 10^{-6}$ when we use Weitzman's parameterisation, or $\approx 8.19 \,\mathrm{x}\, 10^{-5}$ according to our high damage specification, where D_t is assumed to be equal to 0.5 when the atmospheric temperature is 4degC above the pre-industrial level.

Damages are then partitioned between output and capital, or output and TFP, depending on the growth model:

$$D_t^i = f^i \cdot D_t$$

[22] See also Mirrlees and Stern, 1972, who first illustrated this feature in simple optimal growth models.

$$D_t^Y = 1 - \frac{(1 - D_t)}{(1 - D_t^i)}$$

where f is the share of damages to $i = A$ or $i = K$. Integrated assessment models do not in general explicitly address the allocation of damages between capital and output, and vary widely in what they implicitly assume about it. Nordhaus and Boyer's (2000) analysis might be read to suggest that f^K is in the region of 1/3, so 0.3 is the value we choose. The calibration problem is even more acute in the case of allocating damages between output and TFP – there are, as Moyer et al. (forthcoming) also point out, currently severe modelling, data and estimation problems in carrying out such an allocation. Moyer et al. (forthcoming) consequently explore a range of values of f^A between 1% and 100%. We make the relatively conservative assumption that $f^A = 0.05$.

The total abatement cost function is

$$\Lambda_t = \theta_{1,t}\mu_t^{\theta_2}$$

where $\theta_{1,t}$ is a time-varying coefficient and $\theta_2 = 2.8$, hence marginal abatement costs are increasing in emissions control.

Cumulative industrial carbon dioxide emissions are constrained by remaining fossil fuel reserves,

$$\sum_{t=0}^{Tmax} E_t^{IND} \leq CCum$$

where $CCum = 6000$ gigatonnes of carbon is the constraint, and total emissions of carbon are the sum of industrial emissions of carbon dioxide and exogenous emissions of carbon dioxide from land use:

$$E_t = E_t^{IND} + E_t^{LAND}$$

Industrial carbon dioxide emissions at time t are proportional to gross output in the same period, hence there is a different function depending on the growth model:

$$E_t^{IND} = \sigma_t(1 - \mu_t)A_t K_t^{\alpha+\beta} L_t^{1-\alpha}, \text{ or}$$

$$= \sigma_t(1 - \mu_t)\bar{A}_t K_t^{\alpha} L_t^{1-\alpha}$$

where σ_t is the ratio of uncontrolled emissions to output and is an exogenous, time-varying coefficient. It is assumed that $\partial\sigma/\partial t < 0$, representing autonomous improvements in carbon productivity that arise from technical progress and structural change, and that $\partial^2\sigma/\partial t^2 > 0$.

The atmospheric stock of carbon is driven by total emissions, in a system of three equations representing the cycling of carbon between three reservoirs, the atmosphere M^{AT}, a quickly mixing reservoir comprising the upper ocean and parts of the biosphere M^{UP}, and the lower ocean M^{LO}:

$$M_t^{AT} = E_t + \phi_{11} M_{t-1}^{AT} + \phi_{21} M_{t-1}^{UP}$$

$$M_t^{UP} = \phi_{12} M_{t-1}^{AT} + \phi_{22} M_{t-1}^{UP} + \phi_{32} M_{t-1}^{LO}$$

$$M_t^{LO} = \phi_{23} M_{t-1}^{UP} + \phi_{33} M_{t-1}^{LO}$$

Cycling is determined by a set of coefficients ϕ_{jk} that govern the rate of transport from reservoir j to k per unit of time.

The change in the atmospheric stock of carbon from the pre-industrial level determines radiative forcing,

$$F_t = F_{2\times CO_2} \cdot \left(log_2 \frac{M_t^{AT}}{\widehat{M^{AT}}} \right) + F_t^{EX}$$

where $\widehat{M^{AT}}$ is the stock of carbon in the atmosphere before the industrial revolution (i.e. in 1750) and F_t^{EX} is exogenous radiative forcing (capturing among other things the forcing due to greenhouse gases other than carbon dioxide) and is time-dependent. The equation of motion of temperature is given by:

$$T_t = T_{t-1} + \kappa_1 \left[F_t - \frac{F_{2\times CO_2}}{S} \left(T_{t-1} \right) - \kappa_2 \left(T_{t-1} - T_{t-1}^{LO} \right) \right] \tag{5}$$

where T^{LO} is the temperature of the lower oceans and evolves according to:

$$T_t^{LO} = T_{t-1}^{LO} + \kappa_3 \left(T_{t-1} - T_{t-1}^{LO} \right)$$

Annex II: further results

Figure A.1: Baseline atmospheric stock of carbon dioxide, 2005-2205. $S = 3$.

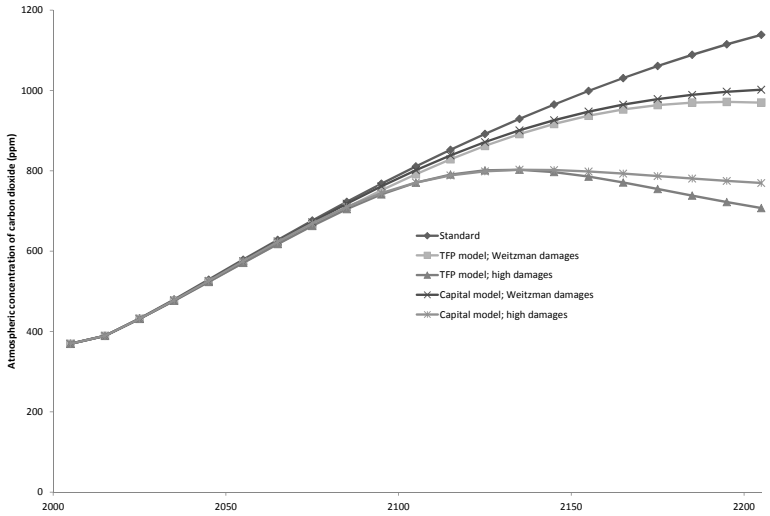

Figure A.2: Baseline global mean temperature (degrees Centigrade above pre-industrial), 2005-2205. The upper panel corresponds with the model of capital damages, while the lower panel corresponds with the model of TFP damages. The damage function calibration is 'Weitzman' unless otherwise indicated.

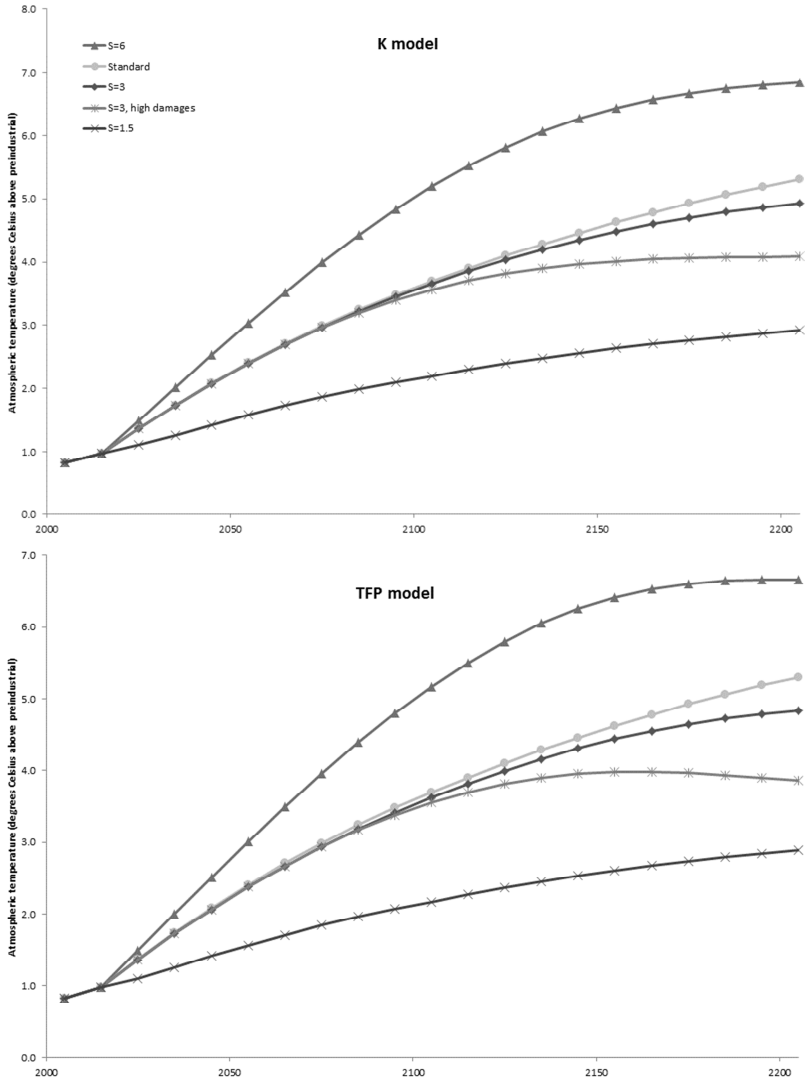

Figure A.3: Optimal atmospheric stock of carbon dioxide, 2005-2205. $S = 3$.

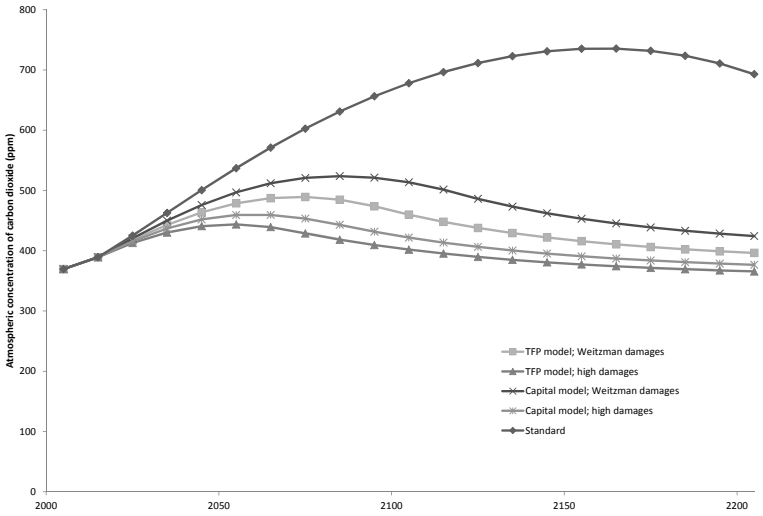

Figure A.4: Optimal global mean temperature (degrees Centigrade above pre-industrial), 2005-2205. The upper panel corresponds with the model of capital damages, while the lower panel corresponds with the model of TFP damages. The damage function calibration is 'Weitzman' unless otherwise indicated.

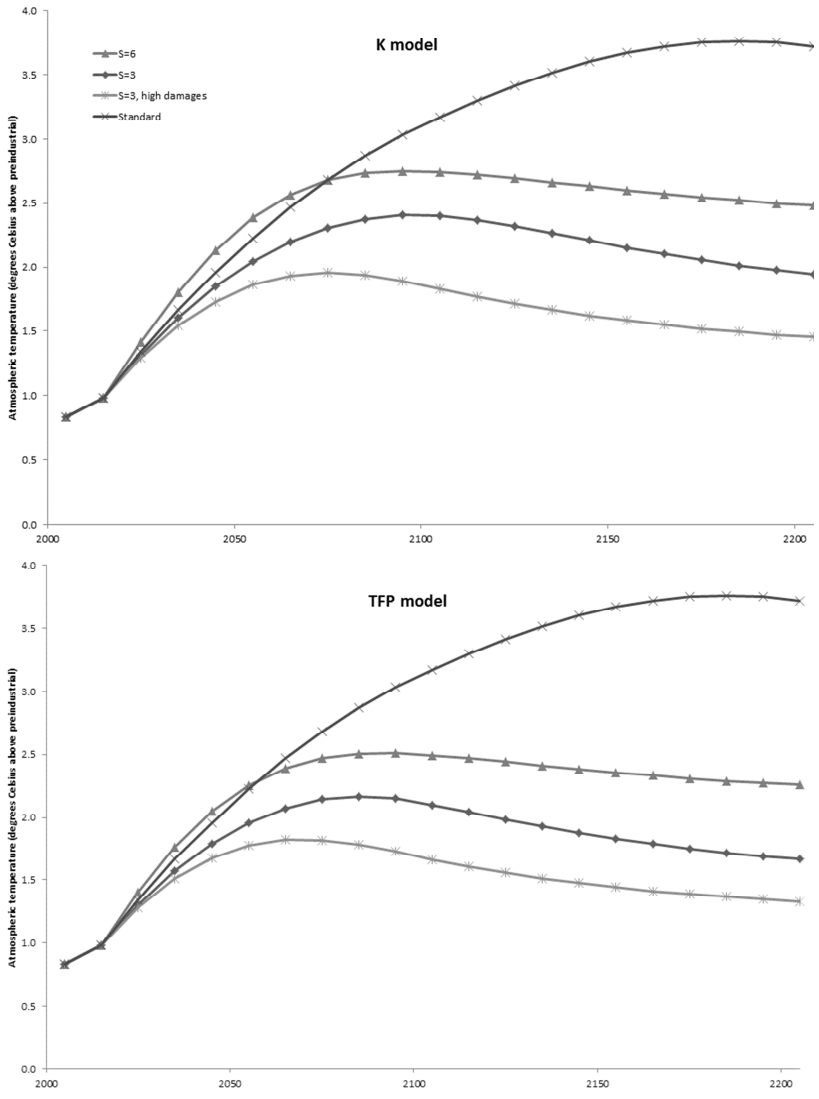

Figure A.5: Optimal atmospheric stock of carbon dioxide, 2005-2205, mean over random S.

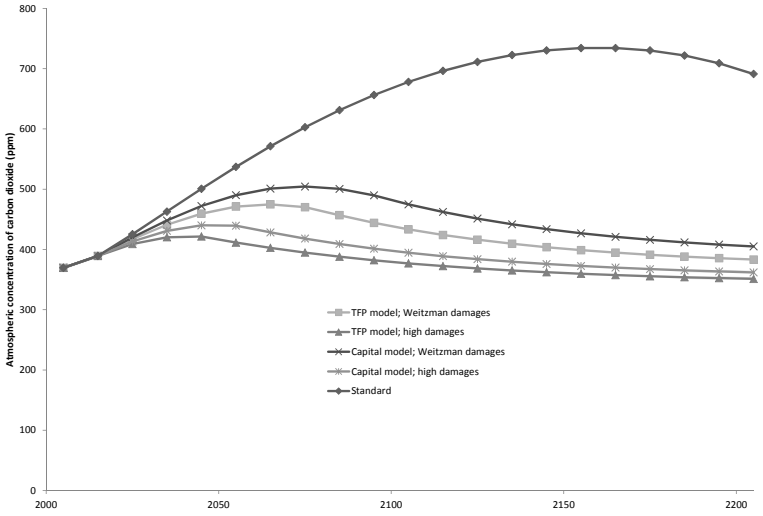

Figure A.6: Optimal global mean temperature (degrees Centigrade above pre-industrial), 2005-2205, mean over random S. The upper panel corresponds with the model of capital damages, while the lower panel corresponds with the model of TFP damages. The bars on the right-hand side give the 90% confidence interval in 2205.