

Centre for
Climate Change
Economics and Policy



Grantham Research Institute on
Climate Change and
the Environment

Climate policy under sustainable discounted utilitarianism

Simon Dietz and Geir B. Asheim

March 2011

**Centre for Climate Change Economics and Policy
Working Paper No. 52**

**Grantham Research Institute on Climate Change and
the Environment**

Working Paper No. 42

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 has five research programmes:

1. Use of climate science in decision-making
2. Mitigation of climate change (including the roles of carbon markets and low-carbon technologies)
3. Impacts of, and adaptation to, climate change, and its effects on development
4. Governance of climate change
5. Management of forests and ecosystems

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

This working paper is intended to stimulate discussion within the research community and among users of research, and its content may have been submitted for publication in academic journals. It has been reviewed by at least one internal referee before publication. The views expressed in this paper represent those of the author(s) and do not necessarily represent those of the host institutions or funders.

Climate policy under sustainable discounted utilitarianism*

Simon Dietz[†] Geir B. Asheim[‡]

March 10, 2011

Abstract

Empirical evaluation of policies to mitigate climate change has been largely confined to the application of discounted utilitarianism (DU). DU is controversial, both due to the conditions through which it is justified and due to its consequences for climate policies, where the discounting of future utility gains from present abatement efforts makes it harder for such measures to justify their present costs. In this paper, we propose sustainable discounted utilitarianism (SDU) as an alternative principle for evaluation of climate policy. Unlike undiscounted utilitarianism, which always assigns zero relative weight to present utility, SDU is an axiomatically based criterion, which departs from DU by assigning zero weight to present utility if and only if it exceeds future welfare. Using the DICE integrated assessment model to run risk analysis, we show that it is possible for future welfare to be below present utility along a ‘business as usual’ development path. Consequently SDU and DU differ, and willingness to pay for emissions reductions is (sometimes significantly) higher under SDU than under DU. Under SDU, stringent schedules of emissions reductions increase social welfare, even if the discount rate is relatively high.

Keywords: Climate change; discounted utilitarianism; intergenerational equity; sustainable development; sustainable discounted utilitarianism

JEL Classification Numbers: D63, D71, Q01, Q54

*Dietz’s research has been supported by the Grantham Foundation for the Protection of the Environment, as well as the Centre for Climate Change Economics and Policy, which is funded by the UK’s Economic and Social Research Council (ESRC) and by Munich Re. Asheim’s research is part of the activities at the center of Equality, Social Organization, and Performance (ESOP) at the Department of Economics at the University of Oslo. ESOP is supported by the Research Council of Norway. We thank Cameron Hepburn for helpful comments. The usual disclaimers apply.

[†]Grantham Research Institute on Climate Change and the Environment, and Department of Geography and Environment, London School of Economics and Political Science (LSE), Houghton Street, London WC2A 2AE, UK (*tel:* +44 (0) 207 955 7589 *fax:* +44 (0) 207 107 5441 *e-mail:* s.dietz@lse.ac.uk).

[‡]Dept. of Economics, University of Oslo, P.O. Box 1095 Blindern, 0317 Oslo, Norway (*tel:* +47 2285 5498 *fax:* +47 2285 5035 *e-mail:* g.b.asheim@econ.uio.no).

1 Introduction

Empirical evaluation of policies to mitigate climate change has been largely confined to the application of discounted utilitarianism (DU). DU means that one stream of consumption is deemed better than another if and only if it generates a higher sum of utilities discounted by a constant per period discount factor δ , where δ is positive and smaller than one.

In spite of its prevalence, DU is controversial, both due to the conditions through which it is justified and due to its consequences for choice in economically relevant situations, such as climate-change policy. As a matter of principle, DU gives less weight to the utility of future generations and therefore treats generations in an unequal manner. If one abstracts from the probability that the world will be coming to an end, thereby assuming that any generation will appear with certainty, it is natural to question whether it is fair to value the utility of future generations less than that of the present one. This criticism has a long tradition dating back at least as far as Ramsey (1928, p. 543), who argued that the practice of discounting later enjoyments in comparison with earlier ones “is ethically indefensible and arises merely from the weakness of the imagination”.

When applied to evaluating climate policies, DU means that the future utility gains of present abatement efforts are discounted, which makes it harder for such measures to justify their present costs. This was one of the earliest findings in the economic literature on climate change (cf. Cline, 1992; Nordhaus, 1991).

One way of treating generations equally is to evaluate policies according to undiscounted utilitarianism, whereby future utilities are summed without being discounted. This alternative was highlighted during the debate following the publication of the Stern Review (2007), which, while committed to DU, applied a utility discount rate of very nearly zero. However, such a criterion (or DU with a near-zero utility discount rate) may contradict our ethical intuitions if used to evaluate all investments, as it is most likely to impose heavy sacrifices on the present generation for the benefit of future generations that are likely to be much better off (Arrow, 1999; Dasgupta, 2007; Mirrlees, 1967; Rawls, 1971). The reason for this weakness of undiscounted utilitarianism is that it assigns zero relative weight to present utility under all circumstances, i.e. even when the present is worse off than the future.

Sustainable discounted utilitarianism (SDU), proposed by Asheim and Mitra (2010), avoids the pitfalls of DU (which is too willing to sacrifice future genera-

tions) and undiscounted utilitarianism (which is too willing to sacrifice the present generation).

SDU departs from DU by placing the additional constraint on social welfare evaluation that no weight be given to present utility if it exceeds future welfare. For example, if the future consequences of climate change entail that present wellbeing exceeds that of the future, then SDU takes into account the future benefits of present mitigation efforts, while ignoring their current costs. Therefore, if there is a non-negligible probability that climate change will undermine future wellbeing, then SDU promotes present action more than DU. However, SDU coincides with DU if the future will for sure be better off than the present in spite of climate change.

If the future will be better off than the present, then additional present sacrifice for the benefit of the future may increase the undiscounted sum of present and future utilities. It also increases the verge between present and future wellbeing, thus making the intergenerational distribution more unequal. Therefore, it seems uncontroversial to allow a trade-off between present and future wellbeing in such circumstances. However, if the future will be worse off than the present, then additional present sacrifice leading to a uniform increase of future wellbeing increases the undiscounted sum of present and future utilities *and* decreases inequality. Hence, such a transfer from the present to the future is desirable both from a utilitarian and egalitarian perspective. This is the ethical underpinning for a condition called “Hammond Equity for the Future”, which gives priority to the future in conflicts where the future is worse off than the present.

“Hammond Equity for the Future” is the key condition in the axiomatic basis for SDU, as investigated by Asheim et al. (2010). SDU also satisfies Chichilnisky’s (1996) “No Dictatorship of the Present”. In contrast, DU is in conflict *both* with “Hammond Equity for the Future”, as it allows a trade-off between present and future wellbeing even when the present is better off than the future, *and* “No Dictatorship of the Present”, as the ranking of DU (on the set of bounded streams) does not depend on what happens beyond some finite future point in time.

Compared to DU, imposing “Hammond Equity for the Future” comes at the cost of (i) removing sensitivity to the interests of the present if the present is better off than the future and (ii) relaxing to the set of non-decreasing streams the property that the trade-off between wellbeing in the first two periods be separable from the remainder of the stream. Regarding (i), there is a large literature, starting with Diamond (1965), which has established a conflict between imposing equity conditions

(like equal treatment and “Hammond Equity for the Future”) on the one hand, and remaining sensitive to the interests of every generation on the other. Asheim et al. (2010, Section 4) present a formal analysis showing how imposing on the set of all streams the separability property of DU mentioned under (ii) is in conflict with equity conditions that respect the interests of future generations.

The present paper proposes SDU as an alternative criterion to DU for the evaluation of climate abatement policies, and it seeks to illustrate that substituting SDU for DU matters for empirical evaluation of such policies. We do so by considering the DICE integrated assessment model, built by William Nordhaus, but where we run risk analysis, including alternative specifications for the important parameters determining the climate sensitivity and damage function. Weitzman (2009, 2010a,b) in particular has raised doubt concerning the climate sensitivity and damage function used in the standard DICE model. Our alternative specifications lead to a non-negligible probability that some generation is better off than its descendants, in which case adopting SDU instead of DU matters for the evaluation, so much so that we are able to show aggressive emissions abatement increases social welfare under SDU, even when the utility discount rate is relatively high. By contrast, we confirm that such abatement policies fail to increase social welfare under DU. We also show that the optimal abatement policy is more stringent under SDU than DU.

Our analysis is ethical in nature, asking what our generation as a collective should do to serve the interests of all generations from an impartial perspective. This is different from taking a strategic perspective, asking what contemporary countries or individuals should do to serve their own interests when such actions influence the future strategic actions of other countries and individuals. Nevertheless, since SDU bridges the gap between DU and undiscounted utilitarianism, we believe it is realistic to suggest that our generation should take its recommendations into account.

The paper is organised as follows. In Section 2 we present formally the concept of SDU. While Asheim and Mitra (2010) have already done so in a deterministic setting and without explicit consideration of population growth, empirical evaluation of climate policies does not permit either of these simplifications, so Section 2 extends SDU to variable population and uncertainty. In Section 3 we present risk analysis with DICE, discussing in particular our choice of climate sensitivity and damage function, before in Section 4 reporting the results from our analysis. As we discuss in the concluding Section 5, the present paper should be considered a first effort in combining recent advances in axiomatic theories of intertemporal social choice

(for a survey, see Asheim, 2010) with empirical evaluation of climate-change policies. Nevertheless, we claim that our analysis is strongly indicative of the importance of broadening the basis of climate-policy evaluation from DU to SDU and beyond.

2 What is Sustainable Discounted Utilitarianism?

In the empirical part of this paper we consider only consumption streams which eventually become constant.¹ This setting simplifies the presentation of SDU, and we refer the reader to Asheim and Mitra (2010) for the more general treatment.

Let $c_t > 0$ denote *consumption* in period t , and let ${}_0\mathbf{c} = (c_0, \dots, c_t, \dots)$ be an infinite stream of consumption, where there exists $T \geq 0$ such that $c_t = c_T$ for all $t \geq T$. A consumption stream ${}_0\mathbf{c}$ is called *egalitarian* if $c_t = c_0$ for all $t \geq 0$.

Utility in a period is derived from consumption in that period alone. The *utility function* U is assumed to be strictly increasing, strictly concave, continuous and continuously differentiable for $c > 0$ with $U'(c) \rightarrow \infty$ as $c \rightarrow 0$. Clearly, any utility function with constant relative inequality aversion satisfies these assumptions.

Let $\delta \in (0, 1)$ denote the *utility discount factor*, and let $\rho > 0$ denote the *utility discount rate*, where the relation between δ and ρ is given by

$$\delta = \frac{1}{1 + \rho}. \quad (1)$$

The theoretical presentation of SDU in this section is facilitated by using the utility discount factor δ , while the numerical results in Section 4 are easier to interpret in terms of the utility discount rate ρ . Keeping in mind eq. (1), this should not create confusion.

In the axiomatic analysis of Asheim and Mitra (2010), time periods correspond to non-overlapping generations assumed to follow each other in sequence. In the empirical analysis of this paper, time periods are shorter, set to ten years (given by the time-step of the DICE model). As long as the discount factor is properly adjusted to reflect a plausible trade-off between present utility and future welfare (straightforward when it is a function of the utility discount rate as above), this choice of period-length does not matter. With overlapping generations, discounting from an ethical perspective between different generations should be differentiated

¹We use a modelling horizon from 2005 to 2395 and assume that consumption remains at the 2395 level thereafter.

from the self-interested discounting that people do within their own lifetimes, and our analysis does reflect the need to do such differentiation.

Given any $\delta \in (0, 1)$, the social welfare function (SWF) w defined by

$$w({}_0\mathbf{c}) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t U(c_t) \quad (2)$$

is the *discounted utilitarian* (DU) SWF. Multiplying by $1 - \delta$ ensures that the utility weights $(1 - \delta, (1 - \delta)\delta, (1 - \delta)\delta^2, \dots, (1 - \delta)\delta^t, \dots)$ sum up to one. Such a normalisation is useful, as it makes the utility of each generation comparable to the welfare of the stream. The DU SWF is well-defined for the set of consumption streams which eventually becomes constant. Furthermore, on this set, it is characterized by

$$w({}_0\mathbf{c}) = (1 - \delta)U(c_0) + \delta w({}_1\mathbf{c}) \quad (\text{w.1})$$

$$w({}_0\mathbf{c}) = U(c_0) \quad \text{if } {}_0\mathbf{c} \text{ is egalitarian.} \quad (\text{w.2})$$

Clearly, (2) implies (w.1) and (w.2). Conversely, for any ${}_0\mathbf{c} = (c_0, \dots, c_t, \dots)$ with $c_t = c_T$ for all $t \geq T$, it follows from (w.2) that

$$w({}_T\mathbf{c}) = U(c_T) = (1 - \delta) \sum_{t=T}^{\infty} \delta^{t-T} U(c_t).$$

Repeated use of (w.1) now yields (2).

The *sustainable discounted utilitarian* (SDU) SWF modifies DU by requiring that an SDU SWF not be sensitive to the interests of the present generation if the present is better off than the future:

$$W({}_0\mathbf{c}) = \begin{cases} (1 - \delta)U(c_0) + \delta W({}_1\mathbf{c}) & \text{if } U(c_0) \leq W({}_1\mathbf{c}) \\ W({}_1\mathbf{c}) & \text{if } U(c_0) > W({}_1\mathbf{c}), \end{cases} \quad (\text{W.1})$$

$$W({}_0\mathbf{c}) = U(c_0) \quad \text{if } {}_0\mathbf{c} \text{ is egalitarian.} \quad (\text{W.2})$$

Condition (W.1) means that future utilities are not discounted (the discount factor is set to 1) if the present is better off than the future. For any ${}_0\mathbf{c} = (c_0, \dots, c_t, \dots)$ with $c_t = c_T$ for all $t \geq T$, it follows from (W.2) that $W({}_T\mathbf{c}) = U(c_T)$. Repeated use of (W.1) now allows us to recursively calculate $W({}_{T-1}\mathbf{c})$, $W({}_{T-2}\mathbf{c})$, and so on, ending up with $W({}_0\mathbf{c})$: Given that $W({}_T\mathbf{c})$, $W({}_{T-1}\mathbf{c})$, \dots , $W({}_t\mathbf{c})$ have been determined, where $1 \leq t \leq T$, the next step in the recursion, $W({}_{t-1}\mathbf{c})$, is determined by

$$W({}_{t-1}\mathbf{c}) = \begin{cases} (1 - \delta)U(c_{t-1}) + \delta W({}_t\mathbf{c}) & \text{if } U(c_{t-1}) \leq W({}_t\mathbf{c}) \\ W({}_t\mathbf{c}) & \text{if } U(c_{t-1}) > W({}_t\mathbf{c}). \end{cases}$$

This explicit algorithm for calculating $W({}_0\mathbf{c})$ is the one used in the empirical analysis. Hence, on our domain of eventually constant consumption streams, the SDU SWF is uniquely determined.

The following result establishes that SDU welfare is a non-decreasing function of time. Furthermore, SDU welfare is bounded above by DU welfare, with the two welfare measures coinciding if the consumption stream is non-decreasing.

Proposition 1 *Assume that ${}_0\mathbf{c}$ is eventually constant.*

(i) *For all $t \geq 0$, $W({}_0\mathbf{c}) \leq W({}_t\mathbf{c}) \leq w({}_t\mathbf{c})$*

(ii) *If ${}_0\mathbf{c}$ is non-decreasing, then $W({}_0\mathbf{c}) = w({}_0\mathbf{c})$.*

Proof. This is a special case of Asheim and Mitra (2010, Proposition 2). ■

Part (ii) means SDU welfare differs from DU welfare only if the consumption stream is not non-decreasing. Hence, existence of some $t \geq 0$ such that $c_t > c_{t+1}$ is a necessary, but insufficient, condition for SDU welfare being strictly below DU welfare, and emphasis will be placed on this possibility in the empirical analysis.

The *stationary equivalent consumption* \bar{c} of a consumption stream ${}_0\mathbf{c}$ is the consumption level \bar{c} , which if held constant yields the same welfare as the consumption stream ${}_0\mathbf{c}$. By (w.2), the stationary equivalent consumption \bar{c} of a consumption stream ${}_0\mathbf{c}$ under DU satisfies $U(\bar{c}) = w({}_0\mathbf{c})$, or since U is strictly increasing:

$$\bar{c} = U^{-1}(w({}_0\mathbf{c})).$$

By (W.2), the stationary equivalent consumption \bar{c} of a consumption stream ${}_0\mathbf{c}$ under SDU satisfies $U(\bar{c}) = W({}_0\mathbf{c})$, or since U is strictly increasing:

$$\bar{c} = U^{-1}(W({}_0\mathbf{c})).$$

We use the stationary equivalent consumption to express non-marginal welfare differences in consumption terms (more on this in Section 4).

2.1 Variable population and uncertainty

Asheim and Mitra (2010) introduce SDU in a deterministic setting where population growth is not explicitly discussed. Application of SDU to climate change, and indeed to a number of other policy issues, requires explicit treatment of population growth and uncertainty, however, and we turn to these issues now.

In Asheim and Mitra (2010, p. 150), consumption in period t is interpreted as “a non-negative indicator of the wellbeing of generation t ”. However, how do we compare the wellbeing of the present generation with the wellbeing of future generations if population size changes over time?

One possibility is to represent the wellbeing of each generation by the product of population size and the utility derived from per-capita consumption. This is the position of ‘classical utilitarianism’. An alternative position is to let the wellbeing of each generation depend only on per-capita consumption; this is called ‘average utilitarianism’. It turns out that applying the condition which provides an axiomatic support for SDU — “Hammond Equity for the Future” — at the level of individuals yields an argument in favour of ‘average utilitarianism’.

At the level of individuals, “Hammond Equity for the Future” is concerned with a situation where all present individuals are equally well off and all future individuals are equally well off, but where the wellbeing of each present individual exceeds that of each future individual. In this situation, “Hammond Equity for the Future”, combined with conditions ensuring representation by means of a SWF, entails that no weight be given to a uniform gain for present individuals. Hence, the comparison is between present and future *individual* wellbeing. Under the assumption that consumption is a comprehensive measure capturing all sources of human wellbeing, this corresponds to a comparison between present and future per-capita consumption, in line with the position of average utilitarianism.

As mentioned in the introduction and further explored by Asheim et al. (2010, Section 3), the appeal of “Hammond Equity for the Future” is that it can be endorsed both from a utilitarian and egalitarian point of view. The utilitarian support rests on there being an infinite number of future individuals. As long as there is an infinite number of generations, each with a positive number of individuals, this is satisfied independently of how population size develops over time. The egalitarian support for the worse-off future individuals in comparison with the better-off present individuals does not depend on their numbers, only their individual wellbeing.²

On this basis, we apply average SDU in our empirical analysis, and, for comparison, we also apply average DU. At the same time, however, we should acknowledge the substantial literature taking the opposite stance (see, e.g. Meade, 1955; Mir-

²The position of average utilitarianism is also supported in Asheim and Bossert’s (2011) axiomatic analysis of sustainability in the case where population is endogenous, so that each population can control (or at least, influence) the size of the population in the next period.

reels, 1967; Dasgupta, 2001; Blackorby et al., 2005), and the extension of SDU, first in terms of its axiomatic foundations and then in application, to cases of variable population would be a valuable future research project.

Now we turn to the analysis of uncertainty. We make the simplifying assumption that the utility function U not only expresses aversion to inequality over time, but also aversion to risk. It is of interest to separate inequality aversion from risk aversion, but this is outside the scope of the present paper. In any case, the identity between aversion to inequality over time and to risk is standard in the empirical literature on climate-policy evaluation, so our results will be easier to compare to previous studies in this way.

Assume that there is a probability distribution over consumption streams.³ By writing $V(u, w) := \min\{(1 - \delta)u + \delta w, w\}$, it follows from (W.1) that $W({}_0\mathbf{c}) = V(U(c_0), W({}_1\mathbf{c}))$. Since V is a concave function of u and w , it follows from Jensen's inequality that

$$E(V(U(\tilde{c}_0), W({}_1\tilde{\mathbf{c}}))) \leq V(E(U(\tilde{c}_0)), E(W({}_1\tilde{\mathbf{c}}))),$$

with strict inequality if both $U(c_0) < W({}_1\mathbf{c})$ and $U(c_0) > W({}_1\mathbf{c})$ are assigned positive probability. This means that dynamic programming methods cannot be used for calculating the expected value of SDU welfare. However, calculating SDU welfare for each realization and then taking expectations yields a correct result. This is what we do in the empirical analysis.

3 Risk analysis with DICE

In order to examine empirically the differences between climate policy evaluation under SDU and DU, we employ the DICE integrated assessment model of the joint climate-economy system, built by William Nordhaus (we adapt the 2007 version of the model, described in Nordhaus, 2008). In brief, DICE couples a standard Ramsey-Cass-Koopmans model of economic growth to a simple model of the climate system.⁴

³In the empirical part this corresponds to the empirical distribution of 1000 random draws of a Latin Hypercube sample.

⁴IPCC (Houghton et al., 1997) coined the term 'simple climate model' to denote models, which specify the atmosphere, surface and deep oceans as one-dimensional, uniformly mixed boxes, which exchange heat and/or CO₂ with each other. By contrast, atmosphere-ocean general circulation models (AOGCMs), the most complex type of climate model, divide the atmosphere and ocean into a detailed three-dimensional grid, with many longitudinal, latitudinal and vertical points.

Output of a composite good is produced using aggregate capital and labour inputs, augmented by exogenous total factor productivity in a Cobb-Douglas production function. However, industrial production is also associated with the emission of carbon dioxide, which is an input to the simple climate model,⁵ resulting in radiative forcing of the atmosphere and an increase in global mean temperature. The climate model couples back to the economy by means of a so-called ‘damage function’, which is a reduced-form polynomial equation associating a change in temperature with a loss in utility, expressed in terms of equivalent output. The damage function in DICE implicitly takes account of the economy and society’s capacity to adapt to climate change, which reduces the amount of output lost for a given increase in global mean temperature, so that the representative agent is left to choose how much to invest in abating CO₂ emissions from production versus how much to invest in the composite capital good for future consumption. The model is globally aggregated and is resolved in decadal time steps from 2005 up to 2395.

DICE is described in full in Nordhaus (2008), and for the sake of brevity we focus our exposition here on those parts of the model we have modified. Since uncertainty is central to climate policy, we select a subset of eight of the most important parameters in DICE, and specify each as random. Table 1 lists these parameters and the form and parameterisation of their probability distribution. In selecting these eight parameters, we have followed the lead of Nordhaus’ (2008) own risk analysis. However, in the case of two parameters, we have chosen an alternative specification. They describe the climate sensitivity and the curvature of the damage function respectively, and we devote special attention to them below.

The first four parameters in Table 1 play a role in determining CO₂ emissions. Of these four parameters, Kelly and Kolstad (2001) showed that growth in total factor productivity and in population are particularly important. The reason is that, in integrated assessment models such as DICE, growth in CO₂ emissions is proportional to growth in global economic output, which in turn is determined in significant measure by productivity growth and by the stock of labour. In addition, where a classical utilitarian SWF is applied, the larger (smaller) is the population when the impacts of climate change occur, the higher (lower) is the social valuation of climate damage (although, as above, this is not a consideration for us). However, while CO₂ emissions are proportional to output, the proportion is usually assumed to decrease over time due to changes in economic structure away from CO₂-intensive

⁵Alongside exogenous emissions of carbon dioxide from land use.

Table 1. Uncertain parameters for simulation of modified DICE-2007.

<i>Parameter</i>	<i>Units</i>	<i>Functional form</i>	<i>Mean</i>	<i>Standard deviation</i>	<i>Source</i>
Initial growth rate of TFP	Per year	Normal	0.0092	0.004	Nordhaus (2008)
Asymptotic global population	Millions	Normal	8600	1892	Nordhaus (2008)
Rate of decarbonisation	Per year	Normal	-0.007	0.002	Nordhaus (2008)
Total resources of fossil fuels	Billion tons of carbon	Normal	6000	1200	Nordhaus (2008)
Price of backstop technology	US\$ per ton of carbon replaced	Normal	1170	468	Nordhaus (2008)
Transfer coefficient in carbon cycle	Per decade	Normal	0.189	0.017	Nordhaus (2008)
Climate sensitivity	°C per doubling of atmospheric CO ₂	Log-normal	1.099*	0.3912*	Weitzman (2009)
Damage function coefficient α_3	Fraction of global output	Normal	0.082	0.028	Own estimate

*In natural logarithm space.

production activities, and to increases in the efficiency of output with respect to CO₂ emissions in a given activity. In DICE, this is achieved by virtue of a variable representing the ratio of emissions/output, which decreases over time as a function of a rate-of-decarbonisation parameter. A further check on industrial CO₂ emissions is provided in the long run by the finite total remaining stock of fossil fuels, which is also treated here as an uncertain parameter.

The fifth uncertain parameter in Table 1 is the price of a so-called ‘backstop’ technology, which in the context of climate-change mitigation is the price of a technology that is capable of completely nullifying CO₂ emissions. In DICE, the backstop is deployed if the control rate on CO₂ emissions reaches 100%, so it is conceptually the marginal cost of the last unit of emissions abatement. Such a technology could most plausibly be a zero-emissions energy technology such as solar or geothermal power.⁶

⁶It could also be a geo-engineering technology such as artificial trees to sequester atmospheric

The backstop price starts very high (mean = US\$1170/tC), but declines over time. Hence it becomes an important determinant of the cost of abatement in the long run.

The sixth and seventh parameters in Table 1 capture important uncertainties in climate science. At a very high level of abstraction, one can distinguish between (i) uncertainties in climate modelling that derive from the cycling of carbon between its various 'sinks' (the atmosphere, the hydrosphere, the biosphere and the lithosphere), which therefore render forecasts of the atmospheric stock of CO₂ for a given pulse of emissions uncertain, and (ii) uncertainties in the relationship between a rising stock of atmospheric CO₂ and temperature. In DICE, the carbon cycle is represented by a system of equations, each containing several parameters. Here, uncertainty about the carbon cycle is captured in a tractable way by focussing on a parameter that determines the proportion of CO₂ in the atmosphere in a particular time period, which dissolves into the upper ocean in the next period.

3.1 Climate sensitivity

Uncertainty about the relationship between atmospheric CO₂ and temperature is captured by a random climate-sensitivity parameter. The climate sensitivity is the increase in global mean temperature, in equilibrium, that results from a doubling in the atmospheric stock of CO₂. In simple climate models, it is critical in determining how fast and how far the planet is forecast to warm in response to emissions. The IPCC's Fourth Assessment Report compiled a number of recent estimates of the climate sensitivity (IPCC, 2007). It concluded that the best estimate of the climate sensitivity is 3°C, that there is a greater than 66% chance of it falling in the range 2–4.5°C (the IPCC's "likely" range), and a less than 10% chance of it being lower than 1.5°C ("very unlikely"). This leaves around a 17% chance that the climate sensitivity exceeds 4.5°C, and indeed a critical feature of all 18 probability density functions of the climate sensitivity compiled by IPCC is that they have a positive skew, with a long tail of high estimates. These tails can be attributed to uncertainty about feedbacks (Roe and Baker, 2007), related for example to clouds and water vapour, and about the cooling effect of aerosols.

In Nordhaus' (2008) risk analysis with DICE, the random climate-sensitivity parameter is normally distributed with a mean of 3°C and a standard deviation of

CO₂, except that DICE has exogenous emissions of CO₂ from land use.

1.1°C. Compared with the evidence compiled by IPCC, however, this distribution may significantly underestimate the probability of very high values. For example, a value of 6.3°C, which is three standard deviations from the mean of Nordhaus’ normal distribution, is assigned a probability of only around 0.1%, whereas several of the pdfs in IPCC (2007) put the corresponding probability at 5–10%. Similarly, in his review of the evidence, Weitzman (2009) considers that there is a 1% chance of the climate sensitivity exceeding 10°C. According to the above normal distribution, this is less likely than a ‘six sigma’ event, so it has a probability of less than $10^{-7}\%$. Therefore we specify the random climate-sensitivity parameter as lognormally distributed. With the parameterisation in Table 1, it can easily be verified that a value of at least 6.3°C, for example, is associated with a 3% probability.

3.2 The damage function

The final parameter in Table 1 is one element of the damage function linking temperature and output-equivalent losses in utility. In recent years, there has been increasing focus in climate-change economics on this critical function (e.g. Weitzman, 2010a), which is unsurprising when one considers that, without the damage function, the accumulation of atmospheric CO₂ has no consequence for social welfare. In many past studies, including those with DICE, the approach has been to specify losses in utility as a quadratic function of global mean temperature:

$$\Omega(T) = \frac{1}{1 + \alpha_1 T_t + \alpha_2 T_t^2}, \quad (3)$$

where Ω , to keep our nomenclature consistent with Nordhaus (2008), is the proportion of output lost at time t , T is the increase in global mean temperature over the pre-industrial level, and α_1 and α_2 are coefficients.

The coefficients α_1 and α_2 are calibrated on the large literature devoted to estimating the cost of climate change in particular sectors of the economy, such as agriculture, energy, and health (summarised in Parry et al., 2007). This literature provides estimates of varying reliability and validity, but it can generally be concluded that the loss in utility for warming of up to about 3°C is relatively well constrained, and is equivalent to a few percent of output. Unfortunately, what the impacts of climate change will be for larger amounts of global warming remains largely in the realm of guesswork (Weitzman, 2009), due to possible non-linearities in the biophysical and socio-economic response to changes in climate variables, as well as possible singularities in the climate system itself (e.g. a collapse in the Antarctic ice sheet,

or a shutdown in the ocean circulation), all of which are very poorly understood at present. This points the spotlight at the functional form for damages.

There has never been any stronger justification for the assumption of quadratic damages than the general supposition of a non-linear relationship, added to the fact that quadratic functions are of a familiar form to economists, with a tractable linear first derivative (i.e. the marginal benefit function of emissions reductions). However, when extrapolated to large temperature levels, the implications of a quadratic function have recently been cast in doubt. Both Ackerman et al. (2010) and Weitzman (2010b) have shown that, with Nordhaus' (2008) calibration of equation (3), 5°C warming results in a loss of utility equivalent to just 6% of output, despite such warming being equivalent to the difference between the present global mean temperature and the temperature at the peak of the last ice age, while it takes around 18°C of warming for losses in utility to exceed the equivalent of 50% of output.

There are various ways to remedy what is increasingly regarded as an implausible forecast. Following equation (3), utility losses can be ramped up by increasing the coefficients α_1 and α_2 , but only at the expense of unrealistically large losses in utility for the initial 3°C warming. Conceptually, much follows from the specification of the utility function. Working with a standard utility function whose sole argument is consumption of the composite good, we can introduce a higher-order term into the damage function to capture greater non-linearity, as Weitzman (2010b) does.⁷ We specify the following function:

$$\Omega(T) = \frac{1}{1 + \alpha_1 T_t + \alpha_2 T_t^2 + \tilde{\alpha}_3 T_t^7}, \quad (4)$$

where $\tilde{\alpha}_3$ is a normally distributed random coefficient with mean and standard deviation reported in Table 1. The remaining coefficients α_1 and α_2 are as in Nordhaus (2008). If $\tilde{\alpha}_3$ takes its mean value, 5°C warming results in a loss of utility equivalent to around 7% of output, while 50% of output is not lost until the global mean temperature is roughly 11°C above the pre-industrial level. Thus the mean value of function (4) remains fairly conservative at high temperatures. However, when $\tilde{\alpha}_3$ is three standard deviations larger than the mean, 5°C warming triggers an output loss of around 25% of output, and 50% of output is lost when warming reaches just 6°C. This is very close to the specification of Weitzman (2010b). Conversely, at three standard deviations below the mean, $\tilde{\alpha}_3$ is small enough that function (4) virtually

⁷The alternative is to specify utility as a function not only of consumption but also of environmental quality directly, for example indexed by global mean warming.

collapses to function (3), so our risk analysis on the damage function can be said to span the approaches taken by Nordhaus on the one hand and Weitzman on the other.

4 Results

As reported in Section 2, a necessary, but insufficient, condition for SDU welfare to be below DU welfare is that there exists some period $t \geq 0$ such that $c_t > c_{t+1}$. This period t has to be smaller than the model's terminal period T , as we impose constant consumption beyond T . Satisfying this condition will in general depend on severe climate-change damage, since the DICE model, in line with other integrated assessment models, predicts strong growth in production in the absence of such damage. For example, when all the coefficients α_i of the damage function are set to zero, so that damages are 'switched off', global mean consumption per capita in DICE is forecast to grow in real terms from US\$6,667 in 2005, the base year, to US\$26,159 in 2105, and onwards to over US\$80,000 in 2205.⁸ Hence the probability that $c_t > c_{t+1}$ for $0 \leq t < T$ may be low, but as long as it is not zero, SDU may lead to a different evaluation of policies to cut CO₂ emissions than will DU. Therefore we begin our analysis of the modelling results by investigating the probability that per-capita consumption is falling at some point over the modelling horizon, conditional on the schedule of emissions cuts pursued.

To begin with, we examine three such climate-change policies. They are, first, 'business as usual', second, a schedule of emissions cuts to limit the atmospheric concentration of CO₂ to twice its pre-industrial level (560 parts per million, hereafter referred to as the 2 CO₂ policy), and, third, a more aggressive schedule of cuts to limit the concentration of CO₂ to only one-and-a-half times its pre-industrial level (420ppm, hereafter referred to as the 1.5 CO₂ policy). The latter two schedules, the abatement schedules, have both been prominent in recent international negotiations about climate policy. Figure 1 presents estimates of the probability that, for each of these three policies, $c_t > c_{t+1}$ with respect to any two successive time periods between 2005 and 2205. To generate these estimates, we take a Latin Hypercube sample of the eight uncertain parameters in DICE, yielding 1000 random draws.

[FIGURE 1 ABOUT HERE]

⁸Using Nordhaus' (2008) standard values for DICE's variables and parameters.

Figure 1 shows that the probability of falling consumption per capita is positive, but small, under all three policies. Indeed, for the coming century, the probability is virtually the same across the three policies, despite the very different emissions control rate under business as usual compared with those under the two abatement policies. This is because, in a very small number of draws, climate damage is so large as to drive consumption growth negative, irrespective of the emissions controls put in place. We can see from Figure 1 that the probability of this is, for example, 0.002 in 2025.⁹ However, what Figure 1 clearly shows is the pay-off to abatement in the 22nd century, when the probability of falling consumption increases significantly under business as usual, while remaining broadly steady under the 2 CO₂ and 1.5 CO₂ abatement policies.

4.1 Welfare evaluation of emissions cuts

Table 2 goes on to examine what these underlying estimates of consumption per capita mean for SDU and DU welfare.

Before we explain the results, a few words are in order about our measure of welfare changes. In computing social welfare according to SDU and DU, we obtain the value of the two abatement policies compared with business as usual in terms of social welfare, measured in utils. We need to express the change in social welfare due to abatement in consumption-equivalent terms, in order to quantify willingness to pay. However, matters are complicated by the very large changes in social welfare we must contemplate as a result of the risk analysis (e.g. in a future contingency where climate damage is severe under business as usual, but can largely be avoided by abatement). We cannot simply normalise the change in social welfare using the (inverse of the) marginal social welfare of a unit of consumption,¹⁰ because the welfare change may not be marginal, so that the first-order approximation of the utility function may be poor (see also Dietz and Hepburn, 2010). Therefore we turn to the concept of the stationary equivalent consumption, a concept which, following Weitzman (1976), is a standard way of representing social welfare in dynamic settings

⁹In fact, Figure 1 shows that aggressive initial emissions abatement along the 1.5 CO₂ policy path actually makes matters worse for a time, as the high initial cost of abatement drives consumption growth negative earlier than under the other two policies in one of the 1000 draws.

¹⁰Whereby $\frac{1}{c^r} \Delta W$ is our welfare change measure in consumption-equivalent terms, where ΔW is the change in social welfare according to either SDU or DU between one of the two abatement policies on the one hand and business as usual on the other.

Table 2. Change in expected stationary equivalent of 2 CO₂ and 1.5 CO₂ policies compared with business as usual, according to SDU and DU.

<i>Abatement policy</i>	<i>SDU</i>			<i>DU</i>		
	5%	<i>Mean</i>	95%	5%	<i>Mean</i>	95%
2 CO ₂	-0.13	4.26	0.71	-0.13	0.16	0.55
1.5 CO ₂	-1.59	3.44	0.59	-1.60	-0.66	0.34

and which we have already discussed in Section 2.¹¹

Table 2 displays our estimates of the stationary equivalent consumption of the 2 CO₂ and 1.5 CO₂ abatement policies, compared with business as usual, according to both SDU and DU. We report the mean estimate, i.e. the expected change in the stationary equivalent, and also indicate the nature of the underlying distribution of the change in the stationary equivalent by reporting both the 5th and 95th percentiles. The utility discount rate ρ in these calculations is 0.02, thus the per-period (i.e. decadal) discount factor is ~ 0.82 , and the coefficient of relative inequality/risk aversion is set to two. For these (and all subsequent) calculations, we use the full modelling horizon from 2005 to 2395.

The table contains our core result, showing that willingness to pay for emissions abatement is significantly larger under SDU than under DU. For the 2 CO₂ policy, the expected increase in the stationary equivalent is 4.26% under SDU, nearly thirtyfold higher than the corresponding estimate of 0.16% under DU. For the 1.5 CO₂ policy, the expected increase is 3.44% under SDU, but -0.66% under DU. Intriguingly, this policy reduces social welfare according to DU, but according to SDU it *increases* it. These results follow directly from the finding, detailed in Figure 1, that consumption is more likely to fall under business as usual than under either of the two abatement policies. SDU places greater value on these policies than DU as a consequence: they are more likely to guarantee sustainability, defined as non-decreasing wellbeing.

What Table 2 also shows is the influence of uncertainty, specifically the small

¹¹We could instead have applied the balanced growth equivalent (BGE) introduced by Mirrlees and Stern (1972). The BGE of a given amount of social welfare is the initial level of consumption per capita, which, if it grows at a constant annual rate over all time, yields the same amount of social welfare. However, as Anthoff and Tol (2009) show, the stationary equivalent consumption gives the same result as the BGE (independently of the choice of growth rate), provided the utility function exhibits constant relative inequality/risk aversion.

number of random draws in which climate damage is severe. This is evident in comparing the expected change in the stationary equivalent with the 95th percentile change under SDU. For both policies, the expected change is in fact greater than the 95th percentile, indicating that a few random draws (less than 5%) have a change in the stationary equivalent so large as to drive the expectation above the 95th percentile. This is one way of showing that concerns about intergenerational equity and concerns about uncertainty are closely linked in the context of climate change.

Figure 2 explores the sensitivity of the expected change in the stationary equivalent as estimated under both SDU and DU to ρ . We examine values for $\rho \in (0, 0.05)$ (corresponding to a range for the decadal discount factor of 1–0.62). It is evident that, in line with the distribution of near-term abatement costs and longer-term benefits, the expected change in the stationary equivalent of the two abatement policies is a decreasing function of ρ , both under SDU and under DU. Indeed, it falls rapidly as ρ is initially increased from 0.

[FIGURE 2 ABOUT HERE]

However, what is more interesting is that the expected change in the stationary equivalent under SDU holds up to a greater extent than under DU, so that the difference between the two evaluation principles grows. When $\rho = 0$, the two principles yield an identical evaluation. The reason for this can readily be seen by comparing (w.1) and (W.1) in Section 2: when the discount factor approaches unity, the SDU algorithm approaches the DU algorithm.¹² However, as ρ increases, the two algorithms can yield different results depending on the probability of falling consumption per capita, and Figure 2 bears this out. For both policies, the expected change in the stationary equivalent falls and eventually becomes negative under DU, but remains positive under SDU. As ρ rises, the far-off future matters less and less under DU, and it is in the far-off future that the benefits of abatement accrue. However, under SDU the far-off future can continue to receive significant weight, if at some point in time future discounted utility is below present utility. We know from Figure 1 that this is the case.

¹²In the limit, as $\rho \rightarrow 0$, or equivalently, $\delta \rightarrow 1$, it follows from (w.1), (w.2), (W.1) and (W.2) that DU and SDU welfare are determined only by the eventual constant part of ${}_0\mathbf{c}$ beyond T , where $c_t = c_T$ for all $t \geq T$. Then both DU and SDU welfare become insensitive to present wellbeing, as $w({}_0\mathbf{c}) = W({}_0\mathbf{c}) = U(c_T)$, illustrating a problematic aspect of undiscounted utilitarianism.

4.2 Optimal policies

Finally, rather than evaluating exogenous policy settings, it is informative to compare the optimal schedule of emissions abatement under SDU and DU. To do this, we set $\rho = 0.02$, and follow Nordhaus (2008) in simultaneously solving a schedule of emissions control rates $\{\mu_t\}$ (where for each t , $\mu_t \in [0, 1]$) that maximises the expectation of SDU and DU respectively. Thus the emissions control rate is a number between 0 and 1, which controls the emissions intensity of output, and a control rate of μ_t results in a fraction $1 - \mu_t$ of output contributing to emissions.

In an integrated assessment model such as DICE, and especially in running risk analysis, solving this optimisation problem is a non-trivial computational challenge. However, we are able to find a solution using a genetic algorithm (Riskoptimizer) and with two modifications to the basic optimisation problem.¹³ Figure 3 presents the schedule of optimal emissions abatement corresponding to SDU and DU. It can be seen, intuitively, that emissions abatement is at least as high under SDU as it is under DU in every time period, and is considerably higher in some, specifically in the latter half of this century and in the next (top panel). The bottom panel also brings out the differences between the two sets of optimal controls, but it further shows that, nevertheless, optimal annual emissions are increasing under SDU and DU for at least the next one hundred years (albeit much less than under business as usual, given the control rates). This is explained by our choice of $\rho = 0.02$, which favours less aggressive strategies of emissions control, all else equal. Setting ρ closer to zero would see the flow of emissions peaking earlier, under DU and especially under SDU.

[FIGURE 3 ABOUT HERE]

¹³First, we only solve μ_t from 2005 to 2245 inclusive, rather than all the way out to 2395. This considerably reduces the scope of the optimisation problem, in return for making little difference to the results, since in the standard version of DICE $\mu_t = 1$, $t > 2245$ (i.e. abatement yields high benefits relative to costs in the far-off future). Our own results also show that $\mu_t \rightarrow 1$ as $t \rightarrow 2245$. Second, we guide the optimisation by imposing the soft constraint that μ_t is non-decreasing everywhere (via an exponential penalty function when μ_t decreases between any two time periods). Otherwise, the algorithm struggled to find a path towards the global maximum. As a soft constraint, the penalty does not enter the welfare evaluation. We were able to verify that the algorithm's best solution satisfied the property of non-decreasingness in μ_t , and that no solution was found which returned higher SDU/DU, where μ_t was decreasing at any point.

5 Concluding remarks

In this paper we have introduced sustainable discounted utilitarianism (SDU) as an alternative criterion to discounted utilitarianism (DU) for the evaluation of climate-abatement policies and we have conducted a risk analysis with the DICE integrated assessment model in order to find out how much the switch matters empirically. To set the stage for this application, we first extended the concept of SDU to variable population and uncertainty (specifically risk). On the back of recent controversies, we also adjusted the climate sensitivity and damage function used in the standard DICE model, as part of our wider risk analysis. The result is that, with our alternative specifications, there is a non-negligible probability that some generation is better off than its descendants due to the impacts of climate change.

In expectations and at an aggregate level, integrated assessment models like DICE assume that the future will be much better than the present, due largely to the assumption of positive growth in total factor productivity. In our empirical analysis we have exploited the possibility that in contingencies where the climate sensitivity is large (so that the increase in atmospheric CO₂ leads to a large rise in temperature) and temperature rise leads to large damages, development of wellbeing may not be monotonically increasing. When such circumstances are assigned positive probability, SDU more than DU promotes present action against climate change, as our analysis has shown. Hence, moving from DU to SDU matters empirically. Furthermore, this result is robust in the sense that the difference between DU and SDU remains substantial even if a high discount rate is applied (we looked at rates up to 5% per annum).

This last observation is particularly significant. In the introduction, we found ourselves agreeing with the arguments of Arrow, Dasgupta, Rawls and others that the use in climate-policy evaluation of undiscounted utilitarianism, or similarly the use of DU with a near-zero discount rate, leads to unappealing transfers of wealth from the present to the future, when applied consistently across the wider set of investment opportunities. Our analysis shows that concern for the wellbeing of future generations might be better taken into account using SDU with a positive discount rate substantially away from zero. What precisely that rate should be, when used alongside SDU, ought to be the focus of a renewed discussion, which is beyond the scope of this paper. In any case, we have shown that, within a relatively broad range, tough emissions abatement schedules will continue to increase social welfare.

At a spatially disaggregated level, climate change may lead to reduced wellbeing (also when compared to present wellbeing) for certain groups, but not for others. Provided that large-scale compensation schemes will not be undertaken, this applies in particular to those living in geographical areas where climate change is likely to be especially severe, and/or where vulnerability is particularly high. One example is likely to be marginal agricultural regions in Africa, another low-lying coastal communities in South and Southeast Asia. At such a disaggregated level, it will matter much more to apply SDU instead of DU, as SDU in effect does not discount the utility loss due to climate change for those groups that are so severely affected.

Therefore, it will be of great interest to apply the SDU criterion (or a similar criterion – *extended rank-discounted utilitarianism* – proposed by Zuber and Asheim, 2010) for evaluating climate change in models where effects are disaggregated on groups, and compare DU to alternative criteria in such a setting. We will turn to this in future work. However, even the present analysis is strongly indicative of the importance of broadening the basis of climate-policy evaluation from DU to SDU and beyond.

Finally, we should comment on the prospect that SDU might actually be applied in policy-making. SDU is the outcome of an explicitly ethical approach to policy evaluation (and within that, an axiomatic approach). As such, one is challenged to adopt an impartial perspective on questions of intergenerational distribution. In reality of course, the present generation – seen as one of a series of non-overlapping generations – enjoys the autonomy to make its own decisions, and the incentive to behave self-interestedly is strong. From a positive point of view, undiscounted utilitarianism can be criticised for this reason: there is plenty of evidence to show that the utility of future generations is discounted at some substantially positive rate, despite ethical objections. However, our conviction is that the case for DU depends in large part on the assumption that the future will be better off than the present for certain, such that if the present generation believed its decisions could leave the future worse off than it, it could be persuaded to revise those decisions. Hence we do believe that the criterion of SDU might influence the present generation in its evaluation of climate policies, and other policies where the sustainability of wellbeing is under threat.

References

- Ackerman, F., E.A. Stanton and R. Bueno (2010), Fat tails, exponents, extreme uncertainty: simulating catastrophe in DICE. *Ecological Economics* **69**(8), 1657–1665.
- Anthoff, D. and R.S.J. Tol (2009), The impact of climate change on the balanced-growth-equivalent: an application of FUND. *Environmental and Resource Economics* **43**, 351–367.
- Arrow, K.J. (1999), Discounting, morality, and gaming. In *Discounting and Intergenerational Equity*. P.R. Portney and J.P. Weyant. Resources for the Future, Washington, DC.
- Asheim, G.B. (2010), Intergenerational equity. *Annual Review of Economics* **2**, 197–222.
- Asheim, G.B. and W. Bossert (2011), Variable population sustainability. Mimeo, Department of Economics, University of Oslo (in preparation).
- Asheim, G.B. and T. Mitra (2010), Sustainability and discounted utilitarianism in models of economic growth. *Mathematical Social Science* **59**, 148–169.
- Asheim, G.B., T. Mitra and B. Tungodden (2010), Sustainable recursive social welfare functions. *Economic Theory*, forthcoming.
- Blackorby, C., W. Bossert and D. Donaldson (2005), *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*. Cambridge University Press, Cambridge, UK.
- Chichilnisky, G. (1996), An axiomatic approach to sustainable development. *Social Choice and Welfare* **13**, 231–257.
- Cline, W.R. (1992), *The Economics of Global Warming*. Institute for International Economics, Washington, DC.
- Dasgupta, P.S. (2001), *Human wellbeing and the Natural Environment*. Oxford University Press, Oxford, UK.
- Dasgupta, P.S. (2007). Commentary: the Stern Review’s economics of climate change. *National Institute Economic Review* **199**, 4–7.

- Diamond, P. (1965). The evaluation of infinite utility streams. *Econometrica* **33**, 170–177.
- Dietz, S. and C. Hepburn (2010), On non-marginal cost-benefit analysis. Centre for Climate Change Economics and Policy Working Paper 20 and Grantham Research Institute on Climate Change and the Environment Working Paper 18. LSE, London.
- Houghton, J.T., L.G. Meira Filho, D.J. Griggs and K. Maskell (eds) (1997), An Introduction to Simple Climate Models used in the IPCC Second Assessment Report. IPCC Technical Paper II. IPCC, Geneva.
- IPCC (2007), Summary for policymakers. In *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. S. Solomon, D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M. Tignor and H.L. Miller. Cambridge University Press, Cambridge, UK and New York, NY.
- Kelly, D.L. and C. Kolstad (2001), Malthus and climate change: betting on a stable population. *Journal of Environmental Economics and Management* **41**, 135–161.
- Meade, J.E. (1955), *Trade and Welfare*. Oxford University Press, Oxford, UK.
- Mirrlees, J.A. (1967), Optimal growth when the technology is changing. *Review of Economic Studies* (Symposium) **34**, 95–124.
- Mirrlees, J.A. and N.H. Stern (1972), Fairly good plans. *Journal of Economic Theory* **4**, 268–288.
- Nordhaus, W.D. (1991), To slow or not to slow: the economics of the greenhouse effect. *Economic Journal* **101**, 920–937.
- Nordhaus, W.D. (2008), *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press, New Haven, CT and London.
- Parry, M.L., O.F. Canziani, J.P. Palutikof, and co-authors (2007), Technical Summary. In *Climate Change 2007: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. M.L. Parry, O.F. Canziani, J.P. Palutikof, P.J. van der Linden and C.E. Hanson. Cambridge University Press, Cambridge, UK, 23–78.

- Pigou, A.C. (1932), *The Economics of Welfare*. MacMillan, London.
- Ramsey, F.P. (1928), A mathematical theory of saving. *The Economic Journal* **38(152)**, 543-559.
- Rawls, J. (1971), *A Theory of Justice*. Harvard University Press, Cambridge, MA.
- Roe, G.H. and M.B. Baker (2007), Why is climate sensitivity so unpredictable. *Science* **318**, 629-632.
- Stern, N. (2007), *The Economics of Climate Change: the Stern Review*. Cambridge University Press, Cambridge, UK.
- Weitzman, M.L. (1976), On the welfare significance of national product in a dynamic economy. *Quarterly Journal of Economics* **90**, pp. 156-162.
- Weitzman, M.L. (2009), On modeling and interpreting the economics of catastrophic climate change. *Review of Economics and Statistics* **91**, 1-19.
- Weitzman, M.L. (2010a), What is the “damages function” for global warming – and what difference might it make? *Climate Change Economics* **1**, 57-69.
- Weitzman, M.L. (2010b), GHG targets as insurance against catastrophic climate damages. Mimeo, Department of Economics, Harvard University.
- Zuber, S. and G.B. Asheim (2010), Justifying social discounting: the rank-discounted utilitarian approach. CESifo Working Paper 3192.




Figure 1: Probability of falling consumption per capita for three emissions abatement policies (business as usual, 2 CO₂ and 1.5 CO₂).





Figure 2: Expected change in stationary equivalent consumption per capita under SDU and DU as a function of the utility discount rate, for the 2 CO₂ (top) and 1.5 CO₂ (bottom) policies.





Figure 3: Optimal emissions under SDU and DU in terms of (top) the emissions control rate and (bottom) annual industrial CO₂ emissions.