Centre for Climate Change Economics and Policy

An ESRC Research Centre

The climate beta

Simon Dietz, Christian Gollier and Louise Kessler

April 2015

Centre for Climate Change Economics and Policy
Working Paper No. 215

Grantham Research Institute on Climate Change and the Environment
Working Paper No. 190

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

Grantham Research Institute on Climate Change and the Environment

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. Its second phase started in 2013 and there are five integrated research themes:

1. Understanding green growth and climate-compatible development
2. Advancing climate finance and investment
3. Evaluating the performance of climate policies
4. Managing climate risks and uncertainties and strengthening climate services
5. Enabling rapid transitions in mitigation and adaptation

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. The Institute is funded by the Grantham Foundation for the Protection of the Environment and the Global Green Growth Institute. It has nine research programmes:

1. Adaptation and development
2. Carbon trading and finance
3. Ecosystems, resources and the natural environment
4. Energy, technology and trade
5. Future generations and social justice
6. Growth and the economy
7. International environmental negotiations
8. Modelling and decision making
9. Private sector adaptation, risk and insurance

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

The climate beta

Simon Dietz*, Christian Gollier† and Louise Kessler*

13th April 2015

Abstract

Reducing emissions of CO_2 today is expected to reduce climate damages in the future. In this paper, we examine the question of whether fighting climate change has the additional advantage of reducing the aggregate risk borne by future generations. This raises the question of the 'climate beta', i.e. the elasticity of climate damages with respect to a change in aggregate consumption. Using the DICE integrated assessment model, we show that the climate beta is positive and close to unity, due above all to the effect of uncertainty about technological progress. In estimating the social cost of carbon, this justifies using a relatively larger rate to discount expected climate damages. On the other hand, expected climate damages are themselves made larger by this effect and overall the NPV of emissions reductions today is increased by the climate beta.

Keywords: beta, CCAPM, climate change, discounting, integrated assessment, mitigation, risk, social cost of carbon

JEL codes: Q54

*ESRC Centre for Climate Change Economics and Policy, Grantham Research Institute on Climate Change and the Environment, and Department of Geography and Environment, London School of Economics and Political Science

†IDEI and LERNA, Toulouse School of Economics

We would like to acknowledge the funding of the UK's Economic and Social Research Council via the ESRC Centre for Climate Change Economics and Policy, as well as the Grantham Foundation for the Protection of the Environment.

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

1 Introduction

Because most of the benefits of mitigating climate change arise in the distant future, the choice of the rate at which these benefits should be discounted is a crucial determinant of our collective willingness to reduce emissions of greenhouse gases. The discount-rate controversy that has emerged in the economic literature over the last two decades shows that there is still substantial disagreement about the choice of this parameter for cost-benefit analysis. One source of controversy comes from the intrinsically uncertain nature of these benefits. It is a tradition in economic theory and finance to adapt the discount rate to the risk profile of the flow of net benefits generated by the policy under scrutiny. The underlying intuition is simple. If a policy tends to raise the collective risk borne by the community of risk-averse stakeholders, this policy should be penalised by increasing the discount rate by a risk premium specific to this policy. On the contrary, if a policy tends to hedge collective risk, this insurance benefit should be acknowledged by reducing the rate at which expected net benefits are discounted, i.e. by adding a negative risk premium to the discount rate.

This simple idea can easily be implemented through the Consumption-based Capital Asset Pricing (CCAPM) theory developed by Lucas (1978). Lucas showed that an investment raises intertemporal social welfare if and only if its Net Present Value (NPV) is positive, where the NPV is obtained by discounting the expected cash flow of the investment at a risk-adjusted rate. This investment-specific discount rate is written as

$$r = r_f + \beta\pi,$$

where r_f is the risk-free rate, π is the systematic risk premium and β is the CCAPM beta of the specific investment under scrutiny. It is defined as the elasticity of the net benefit of the investment with respect to a change in aggregate consumption. This means that a marginal project, whose net benefit is risky but uncorrelated with aggregate consumption, should be discounted at the risk-free rate, because implementing such a project has no effect at the margin on the risk borne by the risk-averse representative agent. A project with a positive (resp. negative) β raises (resp. reduces) collective risk and should be penalised (resp. favoured) by discounting its flow of net benefits at a higher (resp. lower) rate.

The objective of this paper is not to offer a new contribution to the debate about the choice of the risk-free rate, or of the systematic risk premium: there have been many of these in the recent past (see Kolstad et al., 2014, for a recent summary). Rather, the aim of this paper is to discuss the CCAPM beta that should be used to value climate mitigation projects. This 'climate beta' should play an important role in the determination of the social cost of carbon, just as an asset beta is known to be the main determinant of the asset price. Indeed, over the last century in the United States financial markets have exhibited a real risk-free rate and a systematic risk premium of around 1% and 3% respectively. Thus assets whose CCAPM betas are respectively 0 and 2 should be discounted at very different rates of 1% and 7% respectively.

In this paper, we use the DICE model (Nordhaus and Boyer, 2000; Nordhaus, 2008; Nordhaus and Sztorc, 2013) to estimate the climate beta. Before launching into the technical details, however, it is useful to explore the arguments for a small or large climate beta. Sandsmark and Vennemo (2007) were the first to examine this question. In their model, the only stochastic parameter represents the intensity of damages – the loss of GDP – associated with a particular increase in global mean temperature. Given this set-up, large damages are simultaneously associated with low aggregate consumption and a large benefit from mitigating climate change. Hence this model yields a negative correlation between consumption and the climate benefits of mitigation, i.e. a negative climate beta. Weitzman (2013) extends this idea of self-protection via emissions reductions to more recent notions of catastrophic climate change, although he does not work within the conventional CCAPM.

Gollier (2012a) proposes an alternative channel driving the climate beta. Suppose that the only source of uncertainty is exogenous emission-neutral technological progress, which determines economic growth. In this context, under 'business-as-usual', rapid technological progress yields at the same time more consumption, more emissions, a larger concentration of CO_2 in the atmosphere and a larger marginal benefit from fighting climate change, provided the damage function is convex (as is classically assumed). Hence this model yields a positive correlation between consumption and the climate benefits of mitigation, i.e. a positive climate beta.

In this paper, we attempt to encompass these two stories, as well as other possible determinants of beta, and we do so in a dynamic model with investment effects on future consumption. Within a Monte-Carlo simulation of the DICE model, we introduce eight key sources of simultaneous uncertainty about the benefits of climate mitigation and about future consumption, and we measure the climate beta for different maturities of our immediate efforts to reduce emissions. We show that the positive effect on beta of uncertain technological progress, emphasised by Gollier, dominates the negative effect on beta of uncertain climate sensitivity and uncertain damages, the latter story being emphasised by Sandsmark and Vennemo (2007). Put another way, emissions reductions actually increase the aggregate consumption risk borne by future generations. This is in line with Nordhaus (2011), who concluded from Monte-Carlo simulations of the RICE-2011 model that *"those states in which the global temperature increase is particularly high are also ones in which we are on average richer in the future."* The advantage of our work is that, as well as offering a strongly empirically grounded characterisation of key uncertainties in the DICE model, we explicitly compute the climate beta, with the aim of contributing to the debate about the discount rate appropriate for climate-mitigation projects.

In the next section we briefly review beta in the context of Lucas' CCAPM. Section 3 describes how we set up and run the DICE model in order to estimate the climate beta, including an explanation of the effect on beta, in principle, of each of the eight random parameters that we introduce. Section 4 sets out our results and Section 5 concludes.

2 The CCAPM beta

In this section, we derive the standard CCAPM valuation principles as in Lucas (1978). Consider a Lucas-tree economy with a von Neumann-Morgenstern representative agent, whose utility function u is increasing and concave and whose rate of pure preference for the present is δ. Her intertemporal welfare at date 0 is

$$W_0 = \sum_{t=0} e^{-\delta t} E\left[u(c_t)\right], \tag{1}$$

where c_t measures her consumption at date t. Because c_t is uncertain from date 0, it is a random variable. We contemplate an action at date 0, which has the consequence of changing the flow of future consumption to $c_t + \varepsilon B_t$, $t = 0, 1, ...,$ where B_t is potentially random and potentially statistically related to c_t. Because ε is small, the change in intertemporal welfare if the action is implemented can be measured in monetary terms by

$$\mathrm{NPV} = \sum_{t=0} e^{-\delta t} EB_t \frac{u'(c_t)}{u'(c_0)} = \sum_{t=0} e^{-r_t t} EB_t, \tag{2}$$

with

$$r_t = \delta - \frac{1}{t} \ln \frac{EB_t u'(c_t)}{EB_t u'(c_0)}. \tag{3}$$

The right-hand side of equation (2) can be interpreted as the NPV of the action, where, for each maturity t, the expected net benefit EB_t is discounted at a risk-adjusted rate r_t, which is in turn defined by equation (3). In order to simplify equation (3), we make three additional assumptions:

1. For all states of nature, the elasticity of the net conditional benefit at date t with respect to a change in consumption at t is constant, so that there exists $\beta_t \in \mathbb{R}$ such that $E\left[B_t \,|c_t\right] = c_t^{\beta_t}$.

2. Consumption follows a geometric brownian motion with drift μ and volatility σ, so that $x_t = \ln c_t/c_0 \sim N(\mu t, \sigma^2 t)$.

3. The representative agent has constant relative risk aversion γ, so that $u'(c_t) = c_t^{-\gamma}$.

This allows us to rewrite equation (3) as follows:

$$r_t = \delta - \frac{1}{t} \ln \frac{E\left[e^{(\beta_t - \gamma)x_t}\right]}{E\left[e^{\beta_t x_t}\right]}. \tag{4}$$

We now use the well-known property that if $x \sim N(a, b^2)$, then for all $k \in \mathbb{R}$, $E\left[\exp(kx)\right] = \exp\left(ka + 0.5k^2b^2\right)$. Applying this result twice in the above equation implies that

$$r_t = \delta + \left(\beta_t\mu + 0.5\beta_t^2\sigma^2\right) - \left[(\beta_t - \gamma)\mu + 0.5(\beta_t - \gamma)^2\sigma^2\right] = r_f + \beta_t\pi, \tag{5}$$

where the risk-free rate r_f equals

$$r_f = \delta + \gamma\mu - 0.5\gamma^2\sigma^2, \tag{6}$$

and the systematic risk premium equals

$$\pi = \gamma\sigma^2. \tag{7}$$

Observe that both the risk-free rate r_f and the systematic risk premium π have a flat term structure in this framework. However, the risk-adjusted discount rate r_t may have a non-constant term structure, which is homothetic in the term structure of the CCAPM beta of the action, β_t.

In the remainder of the paper, we are interested in estimating the term structure $(\beta_1, \beta_2, ...)$ of the climate beta. This can be done by observing that if $E\left[B_t \,|c_t\right] = c_t^{\beta_t}$, then β_t is nothing other than the regressor of $\ln B_t$ with respect to $\ln c_t$:

$$\ln B_t = \beta_t \ln c_t + \xi_t,$$

where c_t and ξ_t are independent random variables. 1000 draws of the Monte-Carlo simulation of the DICE model generate for each maturity t a series $(\ln B_{it}, \ln c_{it})_{i=1,2,...,1000}$, from which the OLS estimate of $\ln B_t$ on $\ln c_t$ gives us the climate beta associated with that maturity.

Larger beta implies a relatively higher discount rate, but at the same time increased expected benefits of mitigation

Before turning to the modelling proper, it is noteworthy that a large climate beta is not necessarily bad news for those who care about climate change. Although a large climate beta implies a relatively larger discount rate to be applied to climate-mitigation projects, it also raises the expected climate benefit EB_t to be discounted:

$$\begin{aligned} \text{NPV} &= \sum_{t=0} e^{-r_t t} EB_t = \sum_{t=0} e^{-r_t t} c_0^{\beta_t} E\left[e^{\beta_t x_t}\right] = \sum_{t=0} e^{-r_t t} c_0^{\beta_t} e^{\left(\beta_t \mu + 0.5\beta_t^2 \sigma^2\right)t} \\ &= \sum_{t=0} c_0^{\beta_t} \exp\left[(-r_f + \beta_t\left(\mu - \gamma\sigma^2\right) + 0.5\beta_t^2\sigma^2)t\right]. \end{aligned}$$

The NPV of the action is increasing in β_t if β_t is larger than $\gamma - (\mu/\sigma^2)$. Over the last century in the United States, we observed $\mu \approx 2\%$ and $\sigma \approx 4\%$. If we take $\gamma = 2$, this implies that $\gamma - (\mu/\sigma^2) \approx -10.5$. Because most actions yield benefits with an elasticity with respect to a change in aggregate consumption larger than -10.5, we conclude that most investment projects see their NPV increased by an increase in their CCAPM beta. The idea is that the mean growth rate of consumption has been so much larger than its volatility in the past that the effect of a larger beta on the expected benefit is much larger than its effect on the discount rate, thereby generating a positive effect on NPV.

3 Estimating beta with DICE

We now develop estimates of the beta of CO_2 emissions abatement using a modified version of William Nordhaus' well-known DICE model. DICE couples a neoclassical growth model to a simple model of the climate system. Output of a composite good is produced using aggregate capital and labour inputs, augmented by exogenous total factor productivity (TFP). However, production also leads to CO_2 emissions, which are an input to the climate model, resulting in an increase in the atmospheric concentration of CO_2, radiative forcing of the atmosphere and an increase in global mean temperature. The climate model is coupled back to the economy via a 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 into account adaptation to climate change, so the planner is left with the possibility of controlling savings/investment, as usual in a neoclassical growth model, and the price/quantity of CO_2 emissions abatement.

Our analysis is based on the 2013 version of the model, which continues the gradual evolution of the model from previous versions. 'DICE-2013R' is extensively described in Nordhaus and Sztorc (2013), so we will limit our discussion in this section to the modifications we have made. These surround eight parameters in the model, which we randomise for the purpose of estimating betas. These eight random parameters represent key uncertainties at all stages in the climate-policy problem from baseline socio-economic development and associated emissions, through the climate response to emissions, to damages and costs of emissions abatement. Our parameter selection is significantly informed by Nordhaus (2008), in which a similar set of eight parameters was chosen for randomisation based on a review of earlier studies with the model. It is also informed by Dietz and Asheim (2012), who modified Nordhaus' set to take into account scientific evidence on the temperature response to radiative forcing, and to allow for the possibility of steep convexity of the damage function.[1] But we build on both of these previous studies by providing calibrations of the various probability distributions using the latest data.

Table 1 summarises the set of random parameters used in this study, including the data used for calibration. The distributions are assumed independent and each is restricted to be either non-negative or non-positive as appropriate. We implement a CO_2 emissions reduction project by removing one unit of industrial emissions in 2015. This amounts to one gigatonne of CO_2 ($GtCO_2$), and since the atmospheric concentration of CO_2 in 2015

[1]Anderson et al. (2014) is the most comprehensive example of stochastic modelling in the DICE framework, randomising all 51 of the model's parameters as part of a global sensitivity analysis. Their results give reasonable support to our selection: depending on the measure (e.g. social cost of carbon, atmospheric temperature in 2105, etc.), between 3/8 and 5/10 of the parameters, whose uncertainties most affected the value of the measure, are in our set. However, these results do not constitute a definitive basis for selecting a subset of parameters for our study: the problem Anderson et al. (2014) faced is that, for many of the parameters, there are no meaningful data on which a probability distribution might be calibrated. Therefore, to ensure consistency, an arbitrary support of +/- 5, 10 or 20% of the best-guess value was imposed on all parameters, even though data indicate that for some parameters (e.g. climate sensitivity) the support is much wider.

Table 1: Uncertain parameters for simulation of modified DICE-2013R.

Parameter	*Functional form*	*Mean*	*Standard deviation*	*Source*	*Effect on β (likely)*
Initial growth rate of TFP (per year)	Normal	0.0084	0.0059	Maddison project and other sources (see text)	+
Asymptotic global population (millions)	Normal	10854	1368	United Nations (2013)	+
Initial rate of decarbonisation (per year)	Normal	-0.0102	0.0064	IEA (2013)	(+)
Price of back-stop technology in 2050 US$/t$CO_2$(2010 prices)	Log-normal	260	51	Edenhofer et al. (2010)	+
Transfer coefficient in carbon cycle (per decade)	Normal*	0.06835	0.0202	Ciais et al. (2013)	(-)
Climate sensitivity °C per doubling of atmospheric CO_2	Log-logistic**	2.9	1.4	IPCC (2013)	(-)
Damage function coefficient α_2 (% GDP)	Normal	0.0025	0.0006	Tol*** (2009)	(-)
Damage function coefficient α_3 (%GDP)	Normal	0.082	0.028	Dietz and Asheim (2012)	(-)

*Truncated from above at 0.1419. ***Truncated from below at 0.75. ***Including corrigenda published in 2014.

is estimated by DICE to be c. 3167GtCO_2, it may indeed be regarded as a marginal reduction, consistent with the definition of beta given above. We assume that the marginal propensity to save is exogenous and we use Nordhaus' (2013) time series of values, whereby it varies over time, but is always c. $0.23 - 0.24$. We take a Latin Hypercube Sample of the parameter space, which has the advantage of sampling evenly from the domain of each probability distribution, with 1000 draws.

Initial growth rate of TFP As a neoclassical growth model, DICE allocates to TFP that portion of output that cannot be explained by capital and labour inputs at their assumed elasticities (0.3 and 0.7 respectively). It follows (e.g. Barro and Sala-i Martin, 2004) that TFP growth plays a very significant role in determining GDP growth and therefore future consumption and CO_2 emissions (also see Kelly and Kolstad, 2001). As Gollier

(2012a) points out, the effect of variation in TFP growth on beta is positive. Higher TFP growth serves as a positive shock on output and consumption, but this in turn leads to higher emissions, higher total damages from climate change and higher marginal damages, thus higher benefits from emissions abatement.[2]

In line with Nordhaus (2008), we choose to randomise a parameter representing the initial rate of TFP growth. The equation of motion for TFP is

$$A_{t+1} = A_t(1 + g_t^A)$$

where A is TFP and g^A is the growth rate of TFP. In turn,

$$g_t^A = g_0^A \left(1 + \delta^A\right)^{-t}$$

where δ^A is the rate of decline of TFP growth. Since δ^A is several times smaller than g_0^A, uncertainty about the initial growth rate has a lasting impact. To calibrate a probability distribution over g_0^A, we use data on historical TFP growth in the US and UK over the period 1820-2010, compiled from multiple sources[3]. Since DICE is an equilibrium model of long-term growth, we use a rolling 30-year average of annual TFP growth (shorter rolling averages would overstate the potential for fluctuations). A normal distribution fits the data best, with mean and standard deviation as reported in Table 1.

Asymptotic global population Population growth is important in determining the scale of the economy and hence aggregate CO_2 emissions (again see Kelly and Kolstad, 2001). Therefore an increase in population growth has the same qualitative effect on beta as an increase in TFP growth; it increases beta, since the scale effect increases aggregate consumption, emissions, total climate damages and the marginal benefits of mitigation.

In DICE population grows according to the following equation of motion:

$$L_{t+1} = L_t \left(\frac{L_\infty}{L_t}\right)^{g^N}$$

[2]Damages in DICE are expressed as a percentage of output and are strictly less than 100% of output for any finite increase in global mean temperature (see below), so the higher total damages from climate change that would accompany higher output and emissions after a productivity shock will nonetheless be smaller than the direct gain in output from the productivity shock.

[3]Bolt and van Zanden (2013); US Census Bureau; US Bureau of Economic Analysis; Feinstein and Pollard (1988); Matthews et al. (1982). We would like to acknowledge the help of Tom McDermott and Antony Millner in collecting these data, although the resulting estimates are our responsibility.

where L is the population, which converges to the asymptotic global population L_∞ according to the growth rate g^N.

We use the latest global population projections of the United Nations (2013) to calibrate a probability distribution over L_∞. According to these projections, the world population will be at an approximate steady state of 10.85 billion in 2100 on the medium (fertility) variant, within a range of 6.75 billion on the low variant to 16.64 billion on the high variant. This is a non-probabilistic range, which can be set against an emerging – though not uncontested (Lutz et al., 2014) – field of probabilistic population forecasting based on Bayesian methods (Raftery et al., 2012). According to these forecasts, the UN's low and high variants are very unlikely to eventuate (i.e. they are suggested to be well outside the 95% confidence interval: Gerland et al., 2014), because they assume fertility is systematically different to the medium scenario in all countries. Taking this persective into account, we fit a normal distribution to the UN population projections, such that the low variant is three standard deviations away from the mean, with the result that the high variant is even further from the mean.

Initial rate of decarbonisation While growth in CO_2 emissions is proportional to growth in GDP in integrated assessment models such as DICE, the proportion is usually assumed to decrease over time due to changes in economic structure away from carbon-intensive production sectors, and to decreases in the emissions intensity of output in a given sector. These are baseline trends, i.e. achieved without the imposition by a planner of a price/quantity constraint on emissions.

A priori, variation in the rate of decarbonisation has an ambiguous effect on beta. For a given path of output, an increase in the rate of decarbonisation reduces the benefits of mitigation, because it lowers emissions and hence total and marginal climate damages. But the path of output is not given; lower damages increase current income and hence they increase capital investment, future output and consumption, emissions and total damages. There is therefore no doubt that an increase in the rate of decarbonisation increases consumption[4], but what happens to the benefits of mitigation depends in principle on the balance between the negative effect on marginal damages of a reduction in emissions intensity and the positive effect on marginal damages of an expansion in production. However, given that in DICE capital depreciation is 10% per annum while the savings rate is c.

[4]Instantaneous damages are a fraction of current output, and investment is a fraction of output after damages.

0.23-0.24, in practice it might be thought unlikely that the positive effect on marginal damages that goes via investment exceeds the negative, direct effect.

In DICE, 'autonomous' decarbonisation 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:

$$E_t^{IND} = \sigma_t(1 - \mu_t)Y_t \tag{8}$$

where E^{IND} represents industrial CO_2 emissions, μ is the control rate of emissions set by the planner, Y is annual output, and σ is the ratio of uncontrolled emissions to output, given by

$$\sigma_{t+1} = \sigma_t(1 + g_t^\sigma)$$

where $g^\sigma < 0$ is the rate of decline of emissions to output, given by

$$g_t^\sigma = g_0^\sigma \left(1 + \delta^\sigma\right)^t$$

with the initial rate of decline of emissions to output being g_0^σ, subject itself to a rate of decline of $\delta^\sigma < 0$. Similar to TFP, δ^σ is around an order of magnitude smaller than g_0^σ, so the latter is key in driving long-run uncertainty about declining emissions intensity.

To calibrate a distribution over g_0^σ we use data from the International Energy Agency (IEA, 2013), which provides the ratio of global CO_2 emissions from fossil fuels to real global GDP for the period 1971-2011, a period in which planned emissions reductions (i.e. through μ) were trivially small at the global level. Again, we partly smooth annual fluctuations by taking a five-year rolling average. The resulting data are fit best by a normal distribution with mean and standard deviation as reported in Table 1.

Price of the backstop technology While beta is a measure of the correlation of the marginal benefits of emissions abatement with consumption, and therefore abatement costs do not play a direct role in its calculation, they nonetheless play an indirect role, since the emissions scenario on which the mitigation project is undertaken involves non-trivial abatement, even in the baseline that represents business as usual. Variation in abatement costs increases beta: an increase in abatement costs, for a given quantity of abatement, decreases income/consumption, but by decreasing income it also decreases industrial emissions in the long run, due to the same investment effect at play in the case of autonomous decarbonisation. This reduces the benefits of mitigation.

In DICE the total cost of abatement as a percentage of annual GDP, Λ, is determined by

$$\Lambda_t = \theta_{1,t}\mu_t^{\theta_2} \tag{9}$$

where θ_1 and θ_2 are coefficients. The time-path of θ_1 is set so that the marginal cost of abatement at $\mu_t = 1$ is equal to the backstop price at t. Hence randomising the backstop price is a way to introduce uncertainty into abatement costs.

We use the findings of an important recent inter-model comparison study (Edenhofer et al., 2010) to update and characterise uncertainty over the backstop price. Edenhofer et al. (2010) assess the cost of limiting warming to below 2degC in five global energy models. A scenario that stabilises the atmospheric stock of CO_2 at 400ppm requires zero emissions by around 2050, so we can use the models' estimates of marginal abatement costs in 2050 as a measure of the backstop price at that time. Marginal costs range from \$150/tCO_2 to \$500, with an average of \$260, all at today's prices. Since the distribution of cost estimates is asymmetric, we use a log-normal distribution. We set the mean to \$260 and posit that the probability of the lowest and highest estimates is 1/1000. We use a comparable emissions scenario in DICE to retrieve, for each value of the backstop price in 2050, the value of the backstop price in 2010, the initial period.

Transfer coefficient in the carbon cycle There are numerous uncertainties, many of them large, about the behaviour of the climate system in response to carbon emissions (e.g. IPCC, 2013). In the structure of DICE's simple climate model, these can be grouped into (i) uncertainties about the carbon cycle, which render estimates of the atmospheric stock of CO_2 for a given emissions scenario imprecise, and (ii) uncertainties about the relationship between the stock of atmospheric CO_2 and global mean temperature.

The atmospheric stock of carbon in DICE is driven by the sum of industrial emissions from (8) and exogenous emissions from land-use. A system of three equations represents 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}:

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

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

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

where total emissions of CO_2 to the atmosphere are E, and the cycling of CO_2 between the reservoirs is determined by a set of coefficients ϕ_{jk} that govern the rate of transport from reservoir j to k per unit of time. We follow Nordhaus' (2008) uncertainty analysis by randomising ϕ_{12}, the coefficient for the transfer of carbon from M^{AT} to M^{UP}. However, we make use of the latest scientific findings from the IPCC's *Fifth Assessment Report* (Ciais et al., 2013) to calibrate ϕ_{12}. In particular, ϕ_{12} may be calibrated by inspecting evidence on the percentage of a pulse of CO_2 emissions that remains in the atmosphere after 100 years. According to the standard parameterisation of DICE-2013R, this would be c. 36%, but the evidence from multiple climate models collected by Ciais et al. (2013) suggests a mean of 41%, with 54% at +2 standard deviations and 28% at -2 standard deviations. We calibrate ϕ_{12} accordingly, however to ensure the DICE carbon cycle maintains physically consistent behaviour at all values of ϕ_{12}, we must set the lower bound at 31% removed. Table 1 provides details.

Variation in ϕ_{12} also has an ambiguous *a priori* effect on beta. Consider a decrease in ϕ_{12}, which means that more CO_2 emissions remain in the atmosphere. Under these circumstances, if to begin with we take the path of 'potential output' as given, more atmospheric CO_2 means increased total damages, hence consumption is reduced and the marginal benefits of mitigation are increased. This would reduce beta. However, the investment effect means that the path of potential output is not given; reduced income at a particular point in time due to greater damages results in lower investment, which depresses future output. This reduces future consumption too, but because it reduces future CO_2 emissions there is a countervailing, negative effect on the benefits of mitigation. As before, we might expect this countervailing investment effect to be small in comparison with the direct positive effect on the marginal benefits of mitigation.

Climate sensitivity Studies that deploy stochastic versions of DICE have overwhelmingly fixed on the climate sensitivity parameter as a means of rendering uncertain the temperature response to atmospheric CO_2. Climate sensitivity is the increase in global mean temperature, in equilibrium, that results from a doubling in the atmospheric stock of CO_2 from the pre-industrial level. In simple climate models, it is indeed critical in determining how fast and how far the planet is forecast to warm in response to emissions.

Variation in climate sensitivity has an ambiguous – but likely negative – effect on beta, with the causal mechanisms being very similar to those at play in the carbon cycle. Higher climate sensitivity means higher damages, lower consumption and higher benefits of mitigation for given output, but with lower income comes lower investment, lower future output and therefore a counter-balancing negative effect on future emissions that tends to reduce the benefits of mitigation.

The equation of motion of temperature in DICE is given by:

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

where F_{t+1} is radiative forcing, which depends on the atmospheric stock of CO_2, $F_{2\times CO_2}$ is the radiative forcing resulting from a doubling in the atmospheric stock of CO_2 from the pre-industrial level, S is 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.

The latest IPCC report (IPCC, 2013) provides a subjective probability distribution for the climate sensitivity, which is the consensus of the panel's many experts. According to this distribution, S is 'likely' to be 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 1degC, 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. Dietz and Stern (2015) find that a log-logistic function has the appropriate tail shape to fit these data[5] (taking the midpoints of the IPCC ranges), and set the scale and shape parameters of the distribution such that the mean S is 2.9degC, and the standard deviation is 1.4degC. In addition, we truncate the distribution from below at 0.75degC in order to again ensure that the DICE climate model exhibits physically consistent behaviour.

Damage function coefficients α_2 and α_3 Damages are one of the most contestable elements of IAMs (see most recently Pindyck, 2013; Stern, 2013) and, by virtue of its accessibility and simplicity in this regard, DICE has become the common means to give expression to competing views. Much of the debate stems from the inability to constrain a reduced-form damage

[5] That is, the log-logistic function has the lowest root-mean-square error of any distribution fitted.

function at global mean temperature increases of more than 3degC, due to the lack of underlying studies. Antipodes in the literature are given by the traditional quadratic form of Nordhaus (2008; 2013) at one end, and at the other end the damage function with an additional term in Weitzman (2012), which is nearly to the seventh power.

Our damage function takes the following form:

$$D_t = 1/\left(1 + \alpha_1 T_t + \alpha_2 T_t^2 + (\alpha_3 T_t)^7\right)$$

where D is aggregate damages as a percentage of GDP and α_i, $i \in \{1, 2, 3\}$ are coefficients. We specify both α_2 and α_3 as random parameters ($\alpha_1 = 0$ as usual). The former coefficient enables us to capture uncertainty about damages that is represented by the spread of data points provided by the existing literature at warming of between 2 and 3degC. In particular, we use the literature review of Tol (2009) to calibrate α_2, which gives it a mean of 0.0025 and a standard deviation of 0.0006. α_2 is also equivalent to the stochastic parameter in the model proposed by Sandsmark and Vennemo (2007). The coefficient α_3 may be calibrated so as to capture the difference in subjective beliefs of modellers about how substantial damages may be at higher temperatures. We follow Dietz and Asheim (2012) in specifying a normal distribution for α_3 that spans existing suggestions, in that at three standard deviations above the mean total damages approximate Weitzman (2012), while at three standard deviations below the mean they approximately reduce to standard quadratic damages. Further details can again by found in Table 1.

An increase in damages reduces consumption and increases the benefits of mitigation for a given path output gross of climate damages, which decreases beta. However, we must once again be mindful of the investment effect that could reduce future output (gross of climate damages), emissions and therefore benefits of mitigation, so the overall qualitative effect of an increase in damages on beta cannot be determined *a priori*, although we might suppose it to be negative.

4 Results

Using the 1000 draws of the Monte Carlo simulation as the source of variation, we can calculate the instantaneous consumption beta of CO_2 emissions abatement. As a function of time, we can then plot its term structure.

Define the benefits of emissions abatement as its avoided damages, in particular as the difference in consumption with and without removing

1GtCO$_2$. Since the marginal propensity to save is exogenous in our model, the benefits of abatement B are then given by

$$B_t = C_t - C_t^{REF}$$
$$B_t = s\left(1 - D_t\right) Y_t - s\left(1 - D_t^{REF}\right) Y_t^{REF}$$

where C denotes consumption, REF denotes reference outcomes before 1GtCO$_2$ is removed, and s is the marginal propensity to save. Note that output here is net of abatement costs from (9).

Beta is then the covariance between the natural logarithm of reference consumption and the natural logarithm of benefits, divided by the variance of reference consumption:

$$\beta_t = \frac{\mathrm{cov}\left[\ln C_t^{REF}, \ln B_t\right]}{\mathrm{var}\left[\ln C_t^{REF}\right]} \tag{10}$$

The discussion above gives us reason to suppose that, in a dynamic model, the beta of CO$_2$ emissions abatement might depend on the path of growth and emissions. Many of the parameter choices we have already described will impact on this, for instance the initial growth rate of TFP and the initial rate of decarbonisation. But one set of exogenous variables that we must still choose is the set of emissions control rates, $\{\mu_t\}$ in (8). Therefore in Figure 1 we plot the term structure of beta for two different emissions-control scenarios. The first scenario corresponds to the baseline in DICE-2013R, 'business as usual'. According to this scenario, μ_t rises gradually from 4% in 2015 to 14% in 2100 and 54% in 2200. The point made previously about emissions abatement being non-trivial even in the baseline is amply illustrated by these numbers. The second scenario is an example of a path in which emissions reductions are deep: it is the so-called 'Lim2T' scenario from DICE-2013R, in which the planner seeks to limit global warming to no more than 2degC. In Lim2T, μ_t is already 33% in 2015 and it hits the maximum of 100% in 2060.[6]

The headline result is that on both emissions scenarios beta is positive: overall, given the various uncertainties we specify, there is a positive correlation between consumption and the benefits of emissions abatement. The magnitude of beta is quite similar on what are two very different emissions

[6]While the different assumptions we make in this study about, for example, climate sensitivity mean that Lim2T is no longer guaranteed to deliver warming equal to 2degC, for the purpose of estimating beta it is a perfectly good example of a stringent mitigation scenario.

Figure 1: The term structure of β_t for two contrasting emissions scenarios.

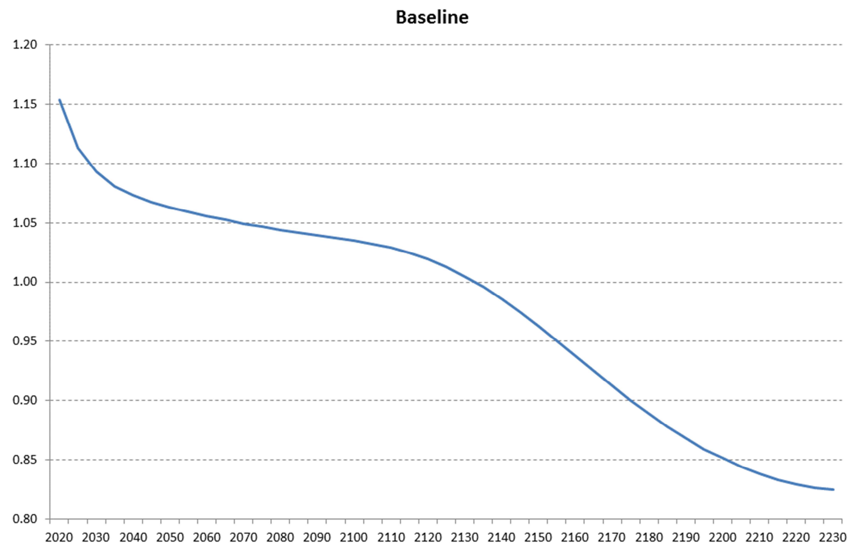

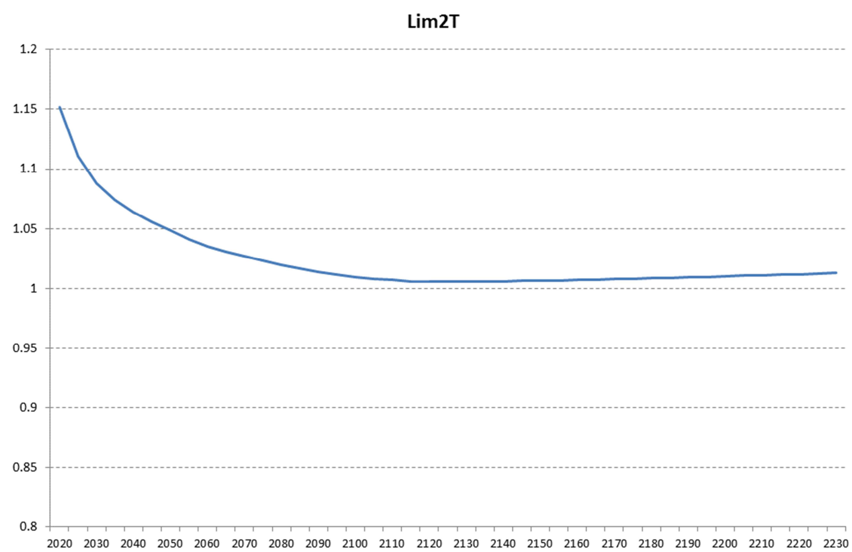

paths, albeit the term structure has a somewhat different profile. If 1GtCO_2 is removed from the baseline, beta starts at 1.15 and falls monotonically but in two distinct stages to 0.83 in 2230. If 1GtCO_2 is removed from Lim2T instead, beta also starts at 1.15, falls to a minimum of just below 1.01 in 2125, before nudging back up fractionally by the end of the horizon.

What is behind these results? We can use two methods of answering this question. First, we can regress the components of beta, i.e. $\ln C_t^{REF}$ and $\ln B_t$, on the full set of uncertain parameters. This should tell us about the relative effects of the different parameters when they vary simultaneously. Second, we can repeat the basic analysis, but focus on each parameter individually. In particular, we hold the parameter in focus to a single value equal to its mean in Table 1, while allowing the other seven parameters to vary according to their distributions. This demonstrates the effect of eliminating uncertainties one by one. The dual of such an analysis would be to look at each random parameter in turn, holding the other seven parameters at a single value, however doing so can, for some parameters, lead to very low variances in $\ln C_t^{REF}$ and unrealistically large absolute values of beta.

The results of our regression analyses can be found in Tables 2 and 3. Table 2 regresses $\ln C_t^{REF}$ on the random parameters for a sample of five time-periods across the modelling horizon, while Table 3 does the same for $\ln B_t$. In both cases, notice that the overall fit of the model is very good. On one level this is unsurprising, since the eight random parameters constitute the only source of variation in the dependent variable. However, it might still have been true that the simple, linear model of main effects that we specify is a poor fit of the data, indicating that second- or higher-order interactions are key. This is not the case.

Looking at the coefficient estimates, where all the parameters have been standardised to aid interpretation, the Tables show all but one of the parameters have the effect on beta that we anticipated. In particular, a one standard-deviation increase in TFP growth has a large, positive and highly statistically significant effect on both $\ln C_t^{REF}$ and $\ln B_t$, thus exerting a large positive effect on beta. An increase in population growth also has a positive and significant effect on $\ln C_t^{REF}$ and $\ln B_t$, but its standardised coefficients are substantially smaller. Working against TFP and population growth, increases in climate sensitivity, α_2 and α_3 have a negative and significant effect on $\ln C_t^{REF}$, while having a positive and significant effect on $\ln B_t$, thus reducing beta. Increasing climate sensitivity has a particularly large effect on $\ln B_t$, but tempering this is the fact that none of these three parameters has an effect on $\ln C_t^{REF}$ that is anything like as substantial as TFP growth. This explains clearly why beta is positive overall. Increasing

Table 2: OLS regression of $\ln\left(C_t^{REF}\right)$ on the set of random parameters.

	2020	*2065*	*2115*	*2165*	*2215*
constant	4.120	4.864	5.298	5.555	5.722
	(0)	(0)	(0.001)	(0.004)	(0.008)
Initial growth rate of TFP	0.062***	0.362***	0.613***	0.767***	0.86***
	(0)	(0)	(0.001)	(0.004)	(0.009)
Asymptotic global population	0.024***	0.096***	0.118***	0.117***	0.112***
	(0)	(0)	(0.001)	(0.004)	(0.008)
Initial rate of decarbonisation	0	-0.002***	-0.013***	-0.043***	-0.083***
	(0)	(0)	(0.001)	(0.004)	(0.008)
Price of back-stop technology in 2050	0	0	0	0.001	0
	(0)	(0)	(0.001)	(0.004)	(0.008)
Transfer coefficient in carbon cycle	0*	0.002***	0.005***	0.013***	0.021***
	(0)	(0)	(0.001)	(0.004)	(0.008)
Climate sensitivity	0***	-0.007***	-0.029***	-0.083***	-0.149***
	(0)	(0)	(0.001)	(0.004)	(0.008)
Damage function coefficient α_2	-0.001***	-0.004***	-0.011***	-0.02***	-0.028***
	(0)	(0)	(0.001)	(0.004)	(0.008)
Damage function coefficient α_3	0	0	-0.004***	-0.027***	-0.056***
	(0)	(0)	(0.001)	(0.004)	(0.008)
R^2	0.999	0.999	0.998	0.969	0.915

Table 3: OLS regression of $\ln(B_t)$ on the set of random parameters.

	2020	*2065*	*2115*	*2165*	*2215*
constant	-8.263	-5.419	-4.801	-4.595	-4.507
	(0.003)	(0.009)	(0.013)	(0.015)	(0.017)
Initial growth rate of TFP	0.062***	0.387***	0.675***	0.868***	0.962***
	(0.003)	(0.011)	(0.015)	(0.017)	(0.019)
Asymptotic global population	0.019***	0.095***	0.121***	0.123***	0.115***
	(0.003)	(0.009)	(0.013)	(0.015)	(0.016)
Initial rate of decarbonisation	0.003	0.038***	0.069***	0.08***	0.053***
	(0.003)	(0.01)	(0.014)	(0.016)	(0.018)
Price of back-stop technology in 2050	0.004	0.009	0.01	0.012	0.015
	(0.003)	(0.009)	(0.013)	(0.015)	(0.016)
Transfer coefficient in carbon cycle	0.006**	-0.087***	-0.139***	-0.136***	-0.102***
	(0.003)	(0.009)	(0.013)	(0.015)	(0.016)
Climate sensitivity	0.09***	0.434***	0.676***	0.795***	0.793***
	(0.003)	(0.009)	(0.013)	(0.015)	(0.017)
Damage function coefficient α_2	0.252***	0.236***	0.206***	0.176***	0.155***
	(0.003)	(0.009)	(0.013)	(0.015)	(0.016)
Damage function coefficient α_3	0.002	0.016*	0.112***	0.198***	0.211***
	(0.003)	(0.009)	(0.013)	(0.015)	(0.016)
R^2	0.901	0.818	0.849	0.86	0.844

the initial rate of decarbonisation and the transfer coefficient in the carbon cycle have statistically significant effects on $\ln C_t^{REF}$ and $\ln B_t$, but they are small in one or both cases. Increasing the transfer coefficient in the carbon cycle reduces beta, while increasing the initial rate of decarbonisation also reduces beta, because it exerts a negative effect on $\ln B_t$ (since the rate of decarbonisation is negative, interpretation of the regression coefficients requires the signs to be reversed). This is the only case in which the dynamic, 'investment' effect outweighs the direct effect. The price of the backstop technology does not have a significant effect on either element.

These analyses also help us explain why beta has a slightly different term structure on the Lim2T emissions scenario than it has on the baseline. On Lim2T the atmospheric concentration of CO_2 is much lower than on the baseline, so the effects of climate sensitivity, α_2 and α_3 on $\ln C_t^{REF}$ and particularly $\ln B_t$ are lower, meaning that the effect of TFP growth comes out still more strongly. Consequently beta does not decline after the beginning of the next century.

Figure 2 comprises a panel of eight charts, each of which plots the term structure of beta when uncertainty about a single parameter is removed. For the sake of brevity, we focus on the baseline scenario. By far the largest difference in the term structure of beta is created when uncertainty about TFP growth is removed. Without it, beta starts at around only 0.6 and falls to a minimum of -2.14 in 2180. This confirms that uncertainty about TFP growth is pivotal in producing an overall positive beta. By contrast, when TFP uncertainty is included, eliminating other uncertainties makes relatively little difference to beta. It is possible only to discern the effect of climate sensitivity on depressing beta later in the modelling horizon (that is, when uncertainty about climate sensitivity is eliminated, beta holds up at around 1.05-1.1, rather than falling to 0.83), and the effect of α_2 on initial values of beta.

Figure 2: The term structure of β_t on the baseline scenario as a function of $N-1$ random parameters.

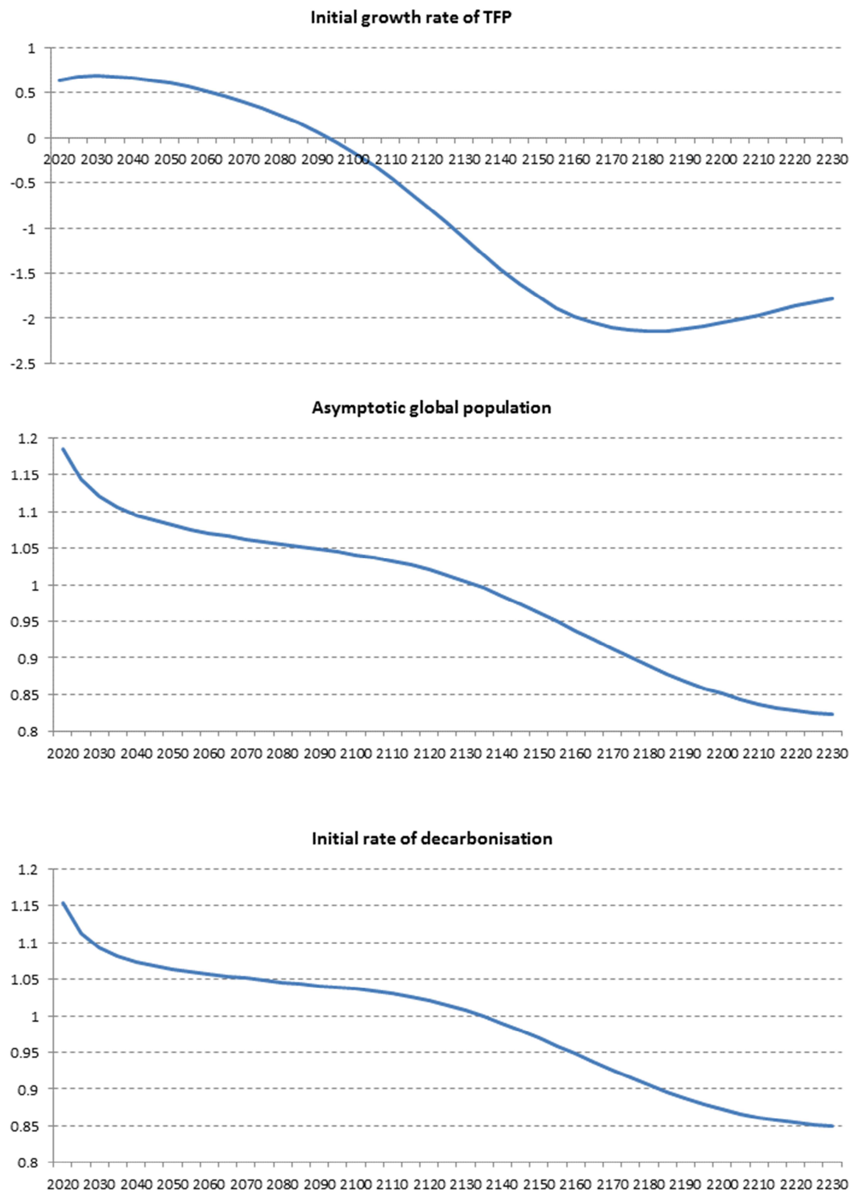

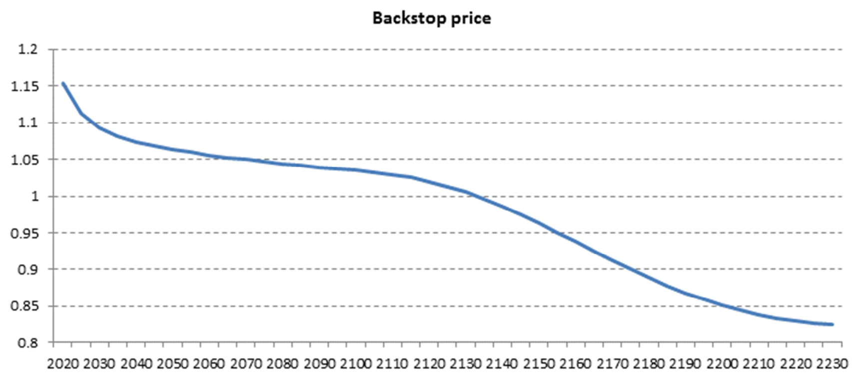
Backstop price
1.2
1.15
1.1
1.05
1
0.95
0.9
0.85
0.8
2020 2030 2040 2050 2060 2070 2080 2090 2100 2110 2120 2130 2140 2150 2160 2170 2180 2190 2200 2210 2220 2230

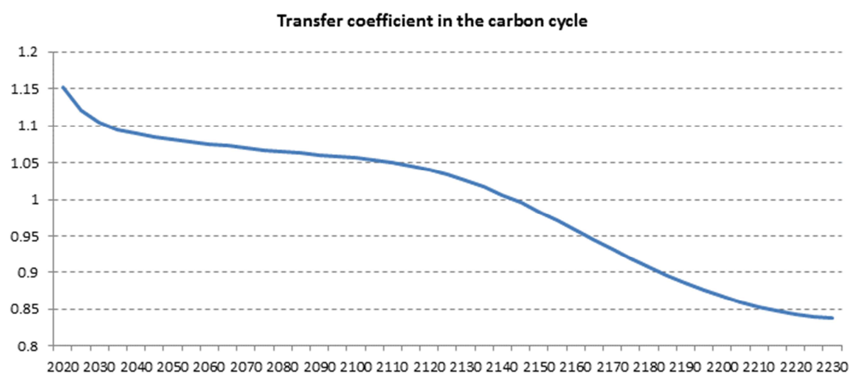
Transfer coefficient in the carbon cycle
1.2
1.15
1.1
1.05
1
0.95
0.9
0.85
0.8
2020 2030 2040 2050 2060 2070 2080 2090 2100 2110 2120 2130 2140 2150 2160 2170 2180 2190 2200 2210 2220 2230

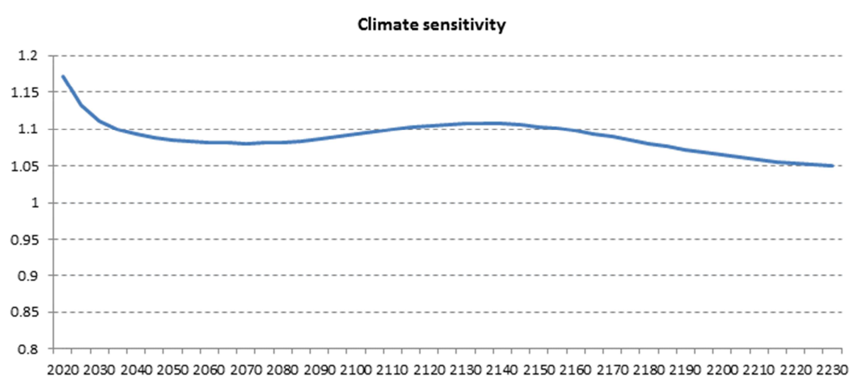
Climate sensitivity
1.2
1.15
1.1
1.05
1
0.95
0.9
0.85
0.8
2020 2030 2040 2050 2060 2070 2080 2090 2100 2110 2120 2130 2140 2150 2160 2170 2180 2190 2200 2210 2220 2230

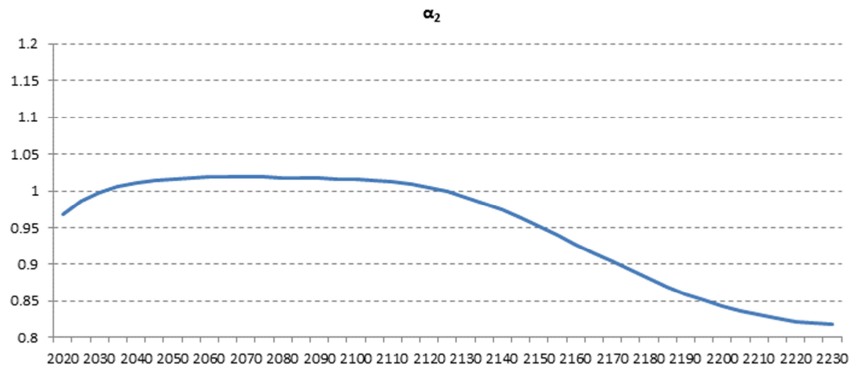
α2
1.2
1.15
1.1
1.05
1
0.95
0.9
0.85
0.8
2020 2030 2040 2050 2060 2070 2080 2090 2100 2110 2120 2130 2140 2150 2160 2170 2180 2190 2200 2210 2220 2230

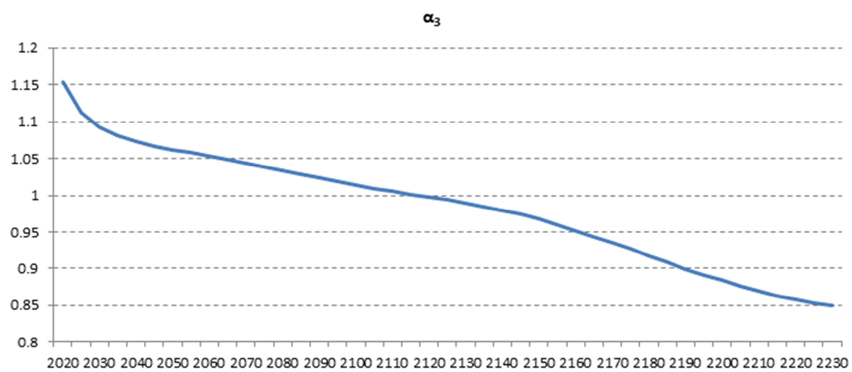
α3
1.2
1.15
1.1
1.05
1
0.95
0.9
0.85
0.8
2020 2030 2040 2050 2060 2070 2080 2090 2100 2110 2120 2130 2140 2150 2160 2170 2180 2190 2200 2210 2220 2230

5 Conclusions

In this paper we have sought to obtain empirically grounded estimates of the climate beta, by bringing together in the DICE model a number of important potential determinants of the covariance of the marginal benefits of climate mitigation with consumption. Naturally the validity of our estimates are, however, affected by the well-known weaknesses shared by all integrated assessment models (e.g. Pindyck, 2013; Stern, 2013).

We find that the climate beta is positive and close to unity throughout the next two centuries and that this holds on two fundamentally different emissions paths, business-as-usual and a path that involves deep cuts with the aim of keeping the global mean temperature below 2degC. The overwhelming driver of these results is uncertainty about technological progress across the whole economy – total factor productivity. Rapid TFP growth is simultaneously associated with higher marginal benefits of emissions reductions and higher consumption. Uncertainty about climate sensitivity and aspects of the damage function provide a countervailing effect that tends to reduce beta, but it is dwarfed by the effect of TFP uncertainty.

Our results patently depend on how TFP uncertainty is calibrated, but they are consistent with previous studies looking at the relative importance of productivity assumptions (Kelly and Kolstad, 2001; Nordhaus, 2011). And, while the structure of DICE assumes multiplicative damages, which contribute a positive, direct relationship between absolute consumption and the absolute benefits of emissions reductions (Weitzman, 2013), it is important to remember that we allow for fat-tailed climate sensitivity and, unlike Nordhaus (2011), for large convexity of the damage function, two of the principal sources of risk of catastrophic climate damages, which contribute to low or negative beta.

Understanding the implications of these findings for climate-change economics requires understanding the dual role played by beta in determining the NPV of mitigation, as set out in Section 2. It is most straightforward to observe that positive beta implies the future benefits of marginal emissions reductions today should be discounted at a relatively higher rate. How much higher?

Two approaches can be followed to answer this question, with radically different conclusions. Both approaches use the CCAPM rule stating that the risk-adjusted discount rate is the sum of terms, the first being a risk-free rate and the second being the product of beta and the systematic risk premium (equations (5), (6) and (7)). The first approach consists in using the systematic risk premium that has been observed in markets, for instance in the United States over the last century, which has been around 3% (see Gollier (2012b), chapter 12). For a project with approximately a unit beta, this means the efficient discount rate for that project should be three percentage points higher than the risk-free rate. The second approach is model-based rather than market-based; one uses the CCAPM formula $\pi = \gamma\sigma^2$ to estimate the risk premium, where σ^2 is the volatility of consumption growth estimated in DICE. According to our simulations, $\sigma^2 = 0.3$ on average over the period 2015-2230, so we obtain a risk premium of only 0.6 percentage points if we accept a coefficient of relative risk aversion $\gamma = 2$, which much of the existing literature would suggest (Kolstad et al., 2014). This leads to a much smaller impact of the positive climate beta on the risk-adjusted climate discount rate.

The large discrepancy between these two recommendations may be explained in part by the fact that our modelling incompletely captures aggregate consumption risk in the real world; we smooth some of the year-to-year volatility in historical productivity growth for the purposes of estimating trend growth (as described in Section 3), and the only novel risk we incorporate is climate change. More generally, however, the discrepancy may be seen as a manifestation of the well-known "equity premium puzzle". Three decades of research on this financial puzzle suggests that the model-based CCAPM approach fails to capture many dimensions of the real world, in particular the existence of structural uncertainties and fat tails (Weitzman, 2007). Although including these dimensions in our model is beyond the reach of this paper – a new concept of beta will need to be developed to accommodate these features – we are inclined to accept this position. We then conclude that a large positive climate beta is important for discounting the future benefits of mitigating climate change.

Is this bad news for those who believe, like us, that climate change should be a primary source of concern for humanity today? Not at all: it is good news, as it will raise the NPV of the future benefits of reducing emissions today. Remember that the climate beta is the elasticity of expected marginal benefits with respect to changes in consumption. Because our modelling suggests consumption growth will be substantial over the next two centuries, the large elasticity that we estimate in this paper also means that the expected marginal benefit will be large. In short, a large beta implies at the same time a larger expected benefit, and a higher rate at which to discount it, with an ambiguous overall effect. However, we have shown in Section 2 in the Gaussian framework that log NPV is quadratic in beta, with a minimum at $\gamma - (\mu/\sigma^2)$, which is negative whether one estimates it with a market- or model-based approach. This implies that the NPV of climate mitigation, i.e. the social cost of carbon, is increasing in beta over its relevant domain. Thus, although it yields a relatively large discount rate for climate-mitigation projects, a large climate beta is good news for those who think that we should do more to reduce carbon emissions.

References

Anderson, B., E. Borgonovo, M. Galeotti, and R. Roson (2014): "Uncertainty in climate change modeling: can global sensitivity analysis be of help?" *Risk Analysis*, 34, 271–293.

Barro, R. J. and X. Sala-i Martin (2004): *Economic Growth*, MIT Press, 2nd ed.

Bolt, J. and J. L. van Zanden (2013): "The first update of the Maddison Project. re-estimating growth before 1820," Maddison-project working paper wp-4, University of Groningen.

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

Ciais, P., C. Sabine, G. Bala, L. Bopp, V. Brovkin, J. Canadell, A. Chhabra, R. DeFries, J. Galloway, M. Heimann, C. Jones, C. L. Quéré, R. Myneni, S. Piao, and P. Thornton (2013): "Carbon and other biogeochemical cycles," in *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, ed. by T. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. Midgley, Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press.

Dietz, S. and G. Asheim (2012): "Climate policy under sustainable discounted utilitarianism," *Journal of Environmental Economics and Management*, 63, 321–335.

Dietz, S. and N. Stern (2015): "Endogenous growth, convexity of damages and climate risk: how Nordhaus' framework supports deep cuts in carbon emissions," *Economic Journal*.

Edenhofer, O., B. Knopf, T. Barker, L. Baumstark, E. Bellevrat, B. Chateau, P. Criqui, M. Isaac, A. Kitous,

S. Kypreos, et al. (2010): "The economics of low stabilization: model comparison of mitigation strategies and costs," *The Energy Journal*, 31, 11–48.

Feinstein, C. H. and S. Pollard (1988): *Studies in Capital Formation in the United Kingdom 1750-1920*, Oxford University Press.

Gerland, P., A. E. Raftery, H. Ševčíková, N. Li, D. Gu, T. Spoorenberg, L. Alkema, B. K. Fosdick, J. Chunn, N. Lalic, et al. (2014): "World population stabilization unlikely this century," *Science*, 346, 234–237.

Gollier, C. (2012a): "Evaluation of long-dated investments under uncertain growth trend, volatility and catastrophes," Tech. rep., Institut d'Économie Industrielle (IDEI), Toulouse.

——— (2012b): *Pricing the Planet's Future: The Economics of Discounting in an Uncertain World*, Princeton University Press.

IEA (2013): "CO2 Emissions from Fuel Combustion 2013 - Highlights," Tech. rep., IEA, Paris.

IPCC (2013): "Working Group I Contribution to the IPCC Fifth Assessment Report: Summary for Policymakers," in *Climate Change 2013: The Physical Science Basis*, IPCC.

Kelly, D. L. and C. D. Kolstad (2001): "Malthus and climate change: betting on a stable population," *Journal of Environmental Economics and Management*, 41, 135–161.

Kolstad, C., K. Urama, J. Broome, A. Bruvoll, M. Cariño Olvera, D. Fullerton, C. Gollier, W. M. Hanemann, R. Hassan, F. Jotzo, M. R. Khan, L. Meyer, and L. Mundaca (2014): "Social, economic and ethical concepts and methods," in *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, ed. by O. Edenhofer, R. Pichs-Madruga, Y. Sokona, E. Farahani, S. Kadner, K. Seyboth, A. Adler, I. Baum, S. Brunner, P. Eickemeier, B. Kriemann, J. Savolainen, S. Schlömer, C. von Stechow, T. Zwickel, and J. Minx, Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press.

Lucas, R. E. (1978): "Asset prices in an exchange economy," *Econometrica*, 46, 1429–1445.

Lutz, W., W. Butz, W. Sanderson, S. Scherbov, et al. (2014): "Population growth: peak probability." *Science*, 346, 561.

Matthews, R. C. O., C. H. Feinstein, and J. C. Odling-Smee (1982): *British Economic Growth, 1856-1973*, Stanford University Press.

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

——— (2011): "Estimates of the social cost of carbon: background and results from the RICE-2011 model," Tech. rep., National Bureau of Economic Research.

——— (2013): *The Climate Casino: Risk, uncertainty, and Economics for a Warming World*, New Haven, CT: Yale University Press.

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

Nordhaus, W. D. and P. Sztorc (2013): "DICE 2013R: Introduction and User's Manual," Tech. rep., Yale University.

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

Raftery, A. E., N. Li, H. Ševčíková, P. Gerland, and G. K. Heilig (2012): "Bayesian probabilistic population projections for all countries," *Proceedings of the National Academy of Sciences*, 109, 13915–13921.

Sandsmark, M. and H. Vennemo (2007): "A portfolio approach to climate investments: CAPM and endogenous risk," *Environmental and Resource Economics*, 37, 681–695.

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, 838–859.

Tol, R. S. (2009): "The economic effects of climate change," *Journal of Economic Perspectives*, 23, 29–51.

United Nations (2013): "World Population Prospects: the 2012 Revision," Tech. rep., United Nations.

WEITZMAN, M. (2012): "GHG targets as insurance against catastrophic climate damages," *Journal of Public Economic Theory*, 14, 221–244.

WEITZMAN, M. L. (2007): "Subjective expectations and asset-return puzzles," *The American Economic Review*, 97, 1102–1130.

——— (2013): "Tail-Hedge Discounting and the Social Cost of Carbon," *Journal of Economic Literature*, 51, 873–82.