The climate beta

Simon Dietz*, Christian Gollier† and Louise Kessler*

4th November 2015

Abstract

Mitigation reduces the expected future damages from climate change, but how does it affect 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. In this paper we show that the climate beta is positive if the main source of uncertainty is exogenous, emissions-neutral technological progress, implying that mitigation has no hedging value. But these results are reversed if the main source of uncertainty is related to the carbon-climate-response and the damage intensity of warming. We then show that in the DICE integrated assessment model the climate beta is positive and close to unity. In estimating the social cost of carbon, this would justify using a relatively high rate to discount expected climate damages. However, the stream of undiscounted expected climate damages is also increasing in the climate beta. We show that this dominates the discounting effect, so that the social cost of carbon is in fact larger than when discounting expected damages at the risk-free rate.

Keywords: beta, 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.

† Toulouse School of Economics.

Email for correspondence: s.dietz@lse.ac.uk. This paper has benefitted greatly from comments by an anonymous referee and participants at EAERE 2016 in Helsinki and the EquiRisk. We would like to acknowledge funding from the UK's Economic and Social Research Council, Amundi, the Sustainable Finance and Responsible Investment chair at IDEI, as well as the Grantham Foundation for the Protection of the Environment.

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 β that should be used to value climate-mitigation projects. This 'climate β' should play an important role in the determination of the social cost of carbon (i.e. the present social value of damages from incremental carbon emissions), just as an asset β is known to be the main determinant of the asset price. Indeed, over the last 150 years in the United States financial markets have exhibited a real risk-free rate and a systematic risk premium of around 1.6% and 4.8 percentage points respectively. Thus assets whose CCAPM betas are respectively 0 and 2 should be discounted at very different rates of 1.6% and 11.2% respectively.[1]

Howarth (2003) was one of the first papers to examine this question. He pointed out that the net benefits of climate-mitigation projects should be discounted at r_f, provided those net benefits are expressed in terms of certainty equivalents (which contains a risk premium). He went on to suggest that the climate β is negative, but did not offer detailed analysis to back up the suggestion.[2] Weitzman's (2007a) *Review of the Stern Review* similarly highlighted the need for further investigation of the correlation between mitigation benefits and consumption, while Aalbers (2009) situated the climate β within a broader set of theoretical conditions, according to which climate-mitigation investments might be discounted at a lower rate than other investments. Sandsmark and Vennemo (2007) provided the first explicit investigation of the climate β. They constructed a simplified climate-economy model, in which 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 climate β. Weitzman (2013) extends the idea that emissions abatement is a hedging strategy against macro-economic risk, in particular invoking potential catastrophic climate change and its avoidance. Most recently, Daniel et al. (2015) have examined a similar problem in the more general context of Epstein-Zin preferences. Their study also suggests a negative climate β, since their estimation of the social cost of carbon is increasing in the degree of risk aversion of the representative agent.

On the other hand, an alternative channel driving the climate β may ex-

[1] See Shiller's dataset: http://www.econ.yale.edu/shiller/data.htm.

[2] In a later paper, Howarth (2009) takes the analysis beyond the scope of the standard CCAPM and therefore of our paper by introducing to the model the concepts of reference dependence and loss aversion. This leads to an additional 'investment risk premium' relating to gains/losses.

ist. Nordhaus (2011) concludes from simulations with 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."* That is, suppose that the only source of uncertainty is exogenous emissions-neutral technological progress, which determines economic growth. In this context, as long as growth is in some measure carbon-intensive, 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). This would yield a positive correlation between consumption and the benefits of mitigation, i.e. a positive climate β. This channel is not present in Sandsmark and Vennemo (2007) and Daniel et al. (2015), because they assume a sure growth rate of pre-climate-damage production and consumption.

In this paper, we attempt to encompass these two stories, as well as other possible determinants of the climate β. We provide two complementary analyses. First, we explore analytical properties of the climate β in a simplified model. As well as serving to develop intuition, the model allows us to explore the role of the structure of climate damages, in particular whether they are multiplicative, as standardly assumed, or additive. We then estimate the climate β numerically using a dynamic integrated assessment model (IAM) 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 β for different maturities of our immediate efforts to reduce emissions.[3] We show that in DICE, which has a multiplicative damage structure, the positive effect on β of uncertain technological progress dominates the negative effect on β of uncertain climate sensitivity and uncertain damages. Put another way, emissions reductions actually increase the aggregate consumption risk borne by future generations. This is in line with Nordhaus (2011), but we go beyond it by explicitly computing the climate β, with the aim of contributing to the debate about the discount rate appropriate for climate-mitigation projects. Another advantage of our work is that we offer a characterisation of key uncertainties in the DICE model that is strongly grounded in underlying data and studies.

In the next section we briefly review β in the context of Lucas' CCAPM. Section 3 describes our analytical model and its results. Section 4 describes

[3] See also Van den Bijgaart et al. (2013) for a recent Monte Carlo simulation of DICE.

how we set up and run the DICE model in order to estimate the climate β numerically. Section 5 sets out our numerical results. In Section 6, we make the important point that a large positive climate β is almost certainly not bad news for those who care about climate change – although it implies a relatively higher discount rate on the benefits of climate-mitigation projects, we identify the conditions under which it also raises the expected (undiscounted) benefits of mitigation. These conditions hold. Section 6 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

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

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 generated by this action is equivalent to an immediate increase in consumption by εNPV, where NPV can be measured as follows:

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

with

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

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 $\mathbb{E}B_t$ 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, which are in line with the classical calibration of the CCAPM model:

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 $\mathbb{E}\left[B_t \,|c_t\right] = c_t^{\beta_t}$.

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:

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

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

Therefore later in the paper we shall be interested in estimating the term structure $(\beta_1, \beta_2, ...)$ of the climate β. This can be done by observing that if $\mathbb{E}\left[B_t \,|c_t\right] = c_t^{\beta_t}$, 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 β associated with that maturity.

3 A simple analytical model of the climate beta

In this section we derive the climate β from a simple analytical model. As well as helping to formalize notions of what determines the climate β, we also use the model to make an important point about the role of the structure of climate damages, specifically what difference it makes to the climate β that damages are multiplicative in most models such as DICE, as opposed to additive.

Let us consider any specific future date t, and let Y represent global economic output within the period $[0, t]$ in the absence of climate damages. Over timescales of decades to centuries, important recent papers in climate science have shown that the increase in the global mean temperature T is approximately linearly proportional to cumulative greenhouse gas emissions (Allen et al., 2009; Matthews et al., 2009; Zickfeld et al., 2009; Goodwin et al., 2015),

$$T = \omega_1 E, \tag{8}$$

where E stands for cumulative industrial greenhouse gas emissions from 0 to t and ω_1 is a parameter called the carbon-climate response (CCR), combining the response of the carbon cycle to emissions and the temperature response to atmospheric carbon (i.e. including radiative forcing and climate sensitivity, etc.). More complex models like DICE deal with these components separately. Emissions are themselves proportional to pre-damage production, so that

$$E = \omega_2 Y - I_0, \tag{9}$$

where $\omega_2 \in [0, 1]$ parameterises the carbon intensity of production, and I_0 is an investment to reduce emissions at the margin.

We assume the damage index D is proportional to increased temperature T at some power k:

$$D = \alpha T^k, \tag{10}$$

where α calibrates the damage function. Parameter k turns out to play an important role in the determination of the climate β in this model. It is widely believed that there is a convex relationship between climate damages and warming, i.e. $k > 1$.

At this stage, let us remain quite general about the way to model the interaction between the damage index D and the index of economic development Y:

$$Q = q(Y, D), \tag{11}$$

where Q is post-damage aggregate output, and q is a bivariate function that is increasing in Y, and decreasing in D, with $Q(Y, 0) = Y$ for all Y. If

$c \in (0,1]$ is the propensity to consume output in period t, then the model yields the following reduced form:

$$C(I_0) = cq\left(Y, \alpha\omega_1^k\left(\omega_2 Y - I_0\right)^k\right). \tag{12}$$

We consider the β of a marginal emissions reduction project. The benefit or cash flow of the project is

$$B \equiv \left.\frac{\partial C}{\partial I_0}\right|_{I_0=0} = -c\alpha\omega_1^k\omega_2^{k-1}Y^{k-1}q_D(Y, hY^k), \tag{13}$$

with $h = \alpha\omega_1^k\omega_2^k$. To sum up, our model characterises the statistical relationship between future consumption $C = C(0)$ and future benefits B as a function of a set of uncertain parameters, such as Y and ω_1. This system is given by the following two equations:

$$\begin{aligned} \ln B &= \ln\left(c\alpha_M\omega_1^k\omega_2^{k-1}\right) + (k-1)\ln Y + \ln\left(-q_D(Y, hY^k)\right), \\ \ln C &= \ln c + \ln q\left(Y, hY^k\right). \end{aligned} \tag{14}$$

How does β respond to the various uncertainties in this model? We proceed one by one through each of the key sources of uncertainty.

3.1 The climate β when the main source of uncertainty is related to exogenous economic growth

Suppose the only source of uncertainty is exogenous, emissions-neutral technological progress, captured in this simplified model by pre-damage production Y. In this case a local estimation of β can be obtained by differentiating the system (14) with respect to Y:

$$\beta \approx \frac{d\ln B/dY}{d\ln C/dY} = \frac{q}{q_D}\frac{(k-1)q_D + Yq_{YD} + Dq_{DD}}{Yq_Y + kDq_D}, \tag{15}$$

where q and its partial derivatives appearing in this equation are evaluated at $\left(Y, hY^k\right)$. The approximation is exact when the uncertainty affecting Y is small.

We calibrate this equation by considering two alternative damage models. In IAMs like DICE, damages are assumed to be multiplicative – proportional to Y – which implies that for instance doubling income also doubles absolute climate damages, all else being equal. We can represent this class of model with the function

$$q(Y, D) = Y(1 - D),$$

Table 1: Calibration of the climate β using equation (16) when the source of uncertainty is exogenous emissions-neutral technological progress.

	$k = 0.5$	$k = 1$	$k = 2$	$k = 3$
$D = 1\%$	0.50	1.01	2.04	3.09
$D = 3\%$	0.51	1.03	2.13	3.31
$D = 5\%$	0.51	1.06	2.24	3.56
$D = 10\%$	0.53	1.13	2.57	4.50
$D = 20\%$	0.57	1.33	4.00	12.00

where D is expressed in percentage points of aggregate income. In this context, (15) simplifies to

$$\beta \approx \frac{k(1-D)}{1-(k+1)D}. \tag{16}$$

In Table 1, we compute the climate β derived from this formula for reasonable values of k and D. It is uniformly positive. Moreover, observe that for damage of less than 5% of GDP,[4] the climate β can be approximated by k. In other words, when the main source of uncertainty is emissions-neutral technological progress, the climate β is approximately equal to the elasticity of climate damage with respect to the increase in global mean temperature. Since the consensus in the damages literature is that $k > 1$, this implies that the climate β is highly likely to be larger than unity, based on this source of uncertainty. What is the intuition behind this result? It is simply that faster technological progress serves as a positive shock to output and consumption, which in turn leads to higher emissions (assuming $\omega_2 > 0$, i.e. provided production is not carbon-free), higher total damages from climate change and higher marginal damages, thus higher benefits from emissions abatement. Future climate benefits of mitigation and future consumption are positively correlated.

Obviously, the fact that damages are assumed to be proportional to pre-damage aggregate income Y plays an important role in this calibration. It is a built-in mechanism towards a positive β. Let us therefore consider an alternative, additive damage structure with

$$q(Y, D) = Y - D,$$

where D measures the absolute level of damages expressed in consumption

[4] The literature on the total economic cost of climate change indicates that it might be at most 5% of GDP when $T = 3$degC (Tol, 2009; IPCC, 2014).

units.[5] In other words, for given warming, doubling pre-damage income has no effect on absolute climate damage. However, the above intuition still applies: increasing income/production results in an increase in emissions as long as $\omega_2 > 0$, which in turn increases temperature and marginal climate damages, if the damage function (10) is convex. So the benefit of mitigation is increased accordingly. What difference then does the additive structure make? When the only source of uncertainty is technological progress via Y,

$$\beta \approx \frac{(k-1)(Y-D)}{Y-kD}. \tag{17}$$

It is interesting to compare Equations (16) and (17), i.e. our estimates of β under multiplicative and additive damages respectively. These two equations are not immediately comparable in fact, because D is expressed in percentage points in the former and in consumption units in the latter. If we express the damage in Eq. (17) in percentage points, $D^{\%} = D/Y$, it can be rewritten as

$$\beta \approx \frac{(k-1)(1-D^{\%})}{1-kD^{\%}}. \tag{18}$$

Eq. (18) is now directly comparable with Eq. (16) and it is clear that the difference lies in replacing k in (16) with $k-1$ in (18). Thus, the numbers in Table 1 also apply in the additive case, except that all betas appearing in this table should be reduced by 1. We summarise this result in the following proposition:

Proposition 1. *Suppose that the main source of uncertainty is emissions-neutral technological progress, and that climate damages are small ($D \leq 5\%$). Then in (a) the multiplicative case, the climate β can be approximated by k, the elasticity of climate damages with respect to warming. In (b) the additive model, the climate β can be approximated by $k-1$.*

When climate damages are large, there is no short-cut to using Equations (16) and (18) in the multiplicative and additive cases respectively to estimate the climate β. Either way, our analysis shows the classical multiplicative model of climate damages has a built-in mechanism towards producing a positive climate β, which is dampened in the additive model. In fact, our analysis shows that there are two independent channels that generate a positive β in the multiplicative case:

[5] The damage function (10) parameter α would need to be recalibrated in order to yield the same absolute damages as in the multiplicative case, for given warming.

- *(convexity effect)* An increase in Y results in higher cumulative emissions E. This in turn increases marginal climate damage – thus the marginal benefit of mitigation – if the damage function (10) is convex, i.e. if $k > 1$;

- *(proportionality effect)* An increase in Y raises damages directly if damages are proportional to Y.

We believe that these two explanations for a positive β in this context have their own merit. The bottom line is that the climate β is positive in this context.

3.2 The climate β when the main source of uncertainty is related to the carbon-climate-response and/or the damage intensity of warming

By contrast, let us now suppose that the only source of uncertainty is the CCR parameter, ω_1. Differentiating the system (14) with respect to ω_1 we obtain

$$\beta \approx \frac{d \ln B / d\omega_1}{d \ln C / d\omega_1} = \frac{q}{q_D} \frac{q_D + D q_{DD}}{D q_D}, \tag{19}$$

where q and its partial derivatives appearing in this equation are again evaluated at (Y, hY^k). The approximation is exact when the uncertainty affecting ω_1 is small. Exactly the same expression for β is obtained when assuming that α rather than ω_1 is uncertain, as examined by Sandsmark and Vennemo (2007) and Daniel et al. (2015). Therefore Equation (19) shows how uncertainty about the CCR and the damage intensity of warming affect the climate β.

Observe that in both the multiplicative and additive models, $q_{DD} = 0$, so that this equation simplifies to

$$\beta \approx \frac{q}{D q_D}, \tag{20}$$

which is unambiguously negative. The intuition for this result is that a higher CCR results in more warming for given cumulative carbon emissions, which in turn yields at the same time higher marginal damage and lower aggregate consumption. Therefore the uncertainty affecting the CCR results in a negative correlation between B and C, and a negative climate β. Similarly, a higher damage intensity of warming as captured by the parameter α results in greater damages for given emissions, and so on.

Proposition 2. *The climate β is unambiguously negative when the main sources of uncertainty are the carbon-climate response and/or the damage intensity of warming.*

This result is independent of whether climate damages are additive or multiplicative in relation to aggregate consumption. For example, in the multiplicative case $q = Y(1 - D)$, the climate β is approximately equal to $-(1 - D)/D$. The same approximation holds in the additive case.[6] If we expect climate damage of around 5% of GDP, we should use a climate β of around -19. There is also an explanation for why the climate β is so large in absolute value in this context. Take the limiting case $\omega_1 = 0$ as a benchmark, and examine the impact of a marginal increase in its value. This will have a marginal (negative) effect on log consumption, but an unbounded effect on the marginal log benefit, since the initial benefit is zero. In other words, fluctuations in ω_1 yield limited relative fluctuations in consumption, but wild relative fluctuations in marginal benefits. This yields a large β in absolute value.

Overall, this analysis illustrates that uncertainty about technological progress on the one hand and about the carbon-climate response and damage intensity of warming on the other hand most likely have contrasting effects on the climate β, the former positive, the latter two negative. This explains the contradictory conclusions that can be found in the literature. Sandsmark and Vennemo (2007) and Daniel et al. (2015) propose models, in which there is no macro-economic uncertainty independent of climate change. Sandsmark and Vennemo (2007) concluded that fighting climate change has a negative CCAPM β. Daniel et al. (2015) corroborate the result of Sandsmark and Vennemo (2007), by showing that the social cost of carbon is increasing in risk aversion in their model. But Nordhaus (2011) contradicts these conclusions by modelling benefits of mitigation that are positively correlated with aggregate consumption. We propose that this contradiction rests in the fact that the Monte-Carlo simulations in Nordhaus (2011) include a source of uncertainty about emissions-neutral technological progress, and it can also be attributed in part to the fact that DICE/RICE deploys a multiplicative damage structure.

[6] Indeed, assuming $q = Y - D$, equation (19) yields $\beta \approx -(Y - D)/D$. This is equal to $-(1 - D^{\%})/D^{\%}$, where $D^{\%} = D/Y$ is the damage expressed as a fraction of Y.

4 Estimating beta with DICE

We now develop estimates of the β 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 multiplicative 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.[7] But we build on both of these previous studies

[7] 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

by providing calibrations of the various probability distributions using the latest data.

Table 2 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 is estimated by DICE to be c. $3167GtCO_2$, it may indeed be regarded as a marginal reduction, consistent with the definition of β 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 discussed in Section 3, the effect of variation in TFP growth on β should be positive.

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. Since we are forecasting more than two centuries into the future and want an appropriately broad range of outcomes for TFP, we want a very long-run series of historical TFP growth, so we use data from

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 2: 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\$/tCO_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.

the US and UK over the period 1820-2010, compiled from multiple sources[8]. 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 2.

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 β as an increase in TFP growth; it increases β, 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:

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

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,

[8]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.

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 β. 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. This is the added dimension of a dynamic model like DICE. So while there is no doubt that an increase in the rate of decarbonisation increases consumption[9], 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. 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:

$$E_t^{IND} = \sigma_t(1 - \mu_t)Y_t \tag{21}$$

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.

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

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 2.

Price of the backstop technology While β 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 β: 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

$$\Lambda_t = \theta_{1,t}\mu_t^{\theta_2} \tag{22}$$

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. Note that together these two uncertainties make up the carbon-climate response in Section 3.

The atmospheric stock of carbon in DICE is driven by the sum of industrial emissions from (21) 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 2 provides details.

Variation in ϕ_{12} also has an ambiguous *a priori* effect on β. 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, then as the simple model from Section 3 showed, more atmospheric CO_2 means increased total damages, hence consumption is reduced and the marginal benefits of mitigation are increased. This would reduce β. 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 β, 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[10] (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 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 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 multiplicative 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

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

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 2.

An increase in damages reduces consumption and increases the benefits of mitigation for a given path output gross of climate damages, which decreases β. 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 β cannot be determined *a priori*, although we might suppose it to be negative.

5 Results

Using the 1000 draws of the Monte Carlo simulation as the source of variation, we can calculate the instantaneous consumption β 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 (22).

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

$$\beta_t = \frac{\text{cov}\left[\ln C_t^{REF}, \ln B_t\right]}{\text{var}\left[\ln C_t^{REF}\right]} \tag{23}$$

The discussion above gives us reason to suppose that, in a dynamic model, the β 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 (21). Therefore in Figure 1 we plot the term structure of β 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.[11]

The headline result is that on both emissions scenarios β 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 β is quite similar on what are two very different emissions paths, albeit the term structure has a somewhat different profile. If 1GtCO_2 is removed from the baseline, β 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, β 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 β, 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 2, 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 β.

The results of our regression analyses can be found in Tables 3 and 4. Table 3 regresses $\ln C_t^{REF}$ on the random parameters for a sample of five

[11] 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 β it is a perfectly good example of a stringent mitigation scenario.

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

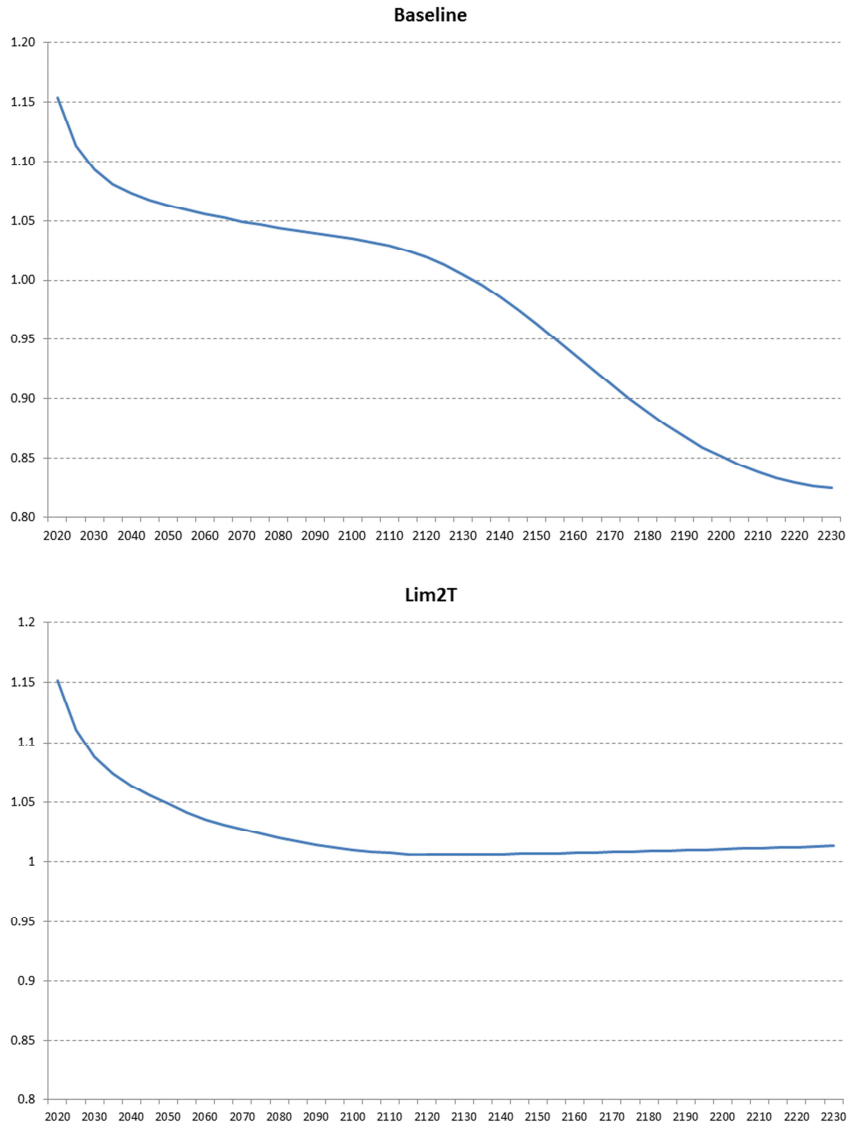

Table 3: OLS regression of $\ln\left(C_t^{REF}\right)$ on the set of random parameters.

	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	0.062***	0.362***	0.613***	0.767***	0.86***
of TFP	(0)	(0)	(0.001)	(0.004)	(0.009)
Asymptotic global	0.024***	0.096***	0.118***	0.117***	0.112***
population	(0)	(0)	(0.001)	(0.004)	(0.008)
Initial rate of	0	-0.002***	-0.013***	-0.043***	-0.083***
decarbonisation	(0)	(0)	(0.001)	(0.004)	(0.008)
Price of back-stop	0	0	0	0.001	0
technology in 2050	(0)	(0)	(0.001)	(0.004)	(0.008)
Transfer coefficient in	0*	0.002***	0.005***	0.013***	0.021***
carbon cycle	(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	-0.001***	-0.004***	-0.011***	-0.02***	-0.028***
coefficient α_2	(0)	(0)	(0.001)	(0.004)	(0.008)
Damage function	0	0	-0.004***	-0.027***	-0.056***
coefficient α_3	(0)	(0)	(0.001)	(0.004)	(0.008)
R^2	0.999	0.999	0.998	0.969	0.915

time-periods across the modelling horizon, while Table 4 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 β 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 β. 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 popula-

Table 4: 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	0.062***	0.387***	0.675***	0.868***	0.962***
of TFP	(0.003)	(0.011)	(0.015)	(0.017)	(0.019)
Asymptotic global	0.019***	0.095***	0.121***	0.123***	0.115***
population	(0.003)	(0.009)	(0.013)	(0.015)	(0.016)
Initial rate of	0.003	0.038***	0.069***	0.08***	0.053***
decarbonisation	(0.003)	(0.01)	(0.014)	(0.016)	(0.018)
Price of back-stop	0.004	0.009	0.01	0.012	0.015
technology in 2050	(0.003)	(0.009)	(0.013)	(0.015)	(0.016)
Transfer coefficient in	0.006**	-0.087***	-0.139***	-0.136***	-0.102***
carbon cycle	(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	0.252***	0.236***	0.206***	0.176***	0.155***
coefficient α_2	(0.003)	(0.009)	(0.013)	(0.015)	(0.016)
Damage function	0.002	0.016*	0.112***	0.198***	0.211***
coefficient α_3	(0.003)	(0.009)	(0.013)	(0.015)	(0.016)
R^2	0.901	0.818	0.849	0.86	0.844

tion 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 β. 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 β is positive overall. Increasing 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 β, while increasing the initial rate of decarbonisation also reduces β, 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 β 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 β 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 β 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 β is created when uncertainty about TFP growth is removed. Without it, β 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 β. By contrast, when TFP uncertainty is included, eliminating other uncertainties makes relatively little difference to β. It is possible only to discern the effect of climate sensitivity on depressing β later in the modelling horizon (that is, when uncertainty about climate sensitivity is eliminated, β holds up at around 1.05-1.1, rather than falling to 0.83), and the effect of α_2 on initial values of β.

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

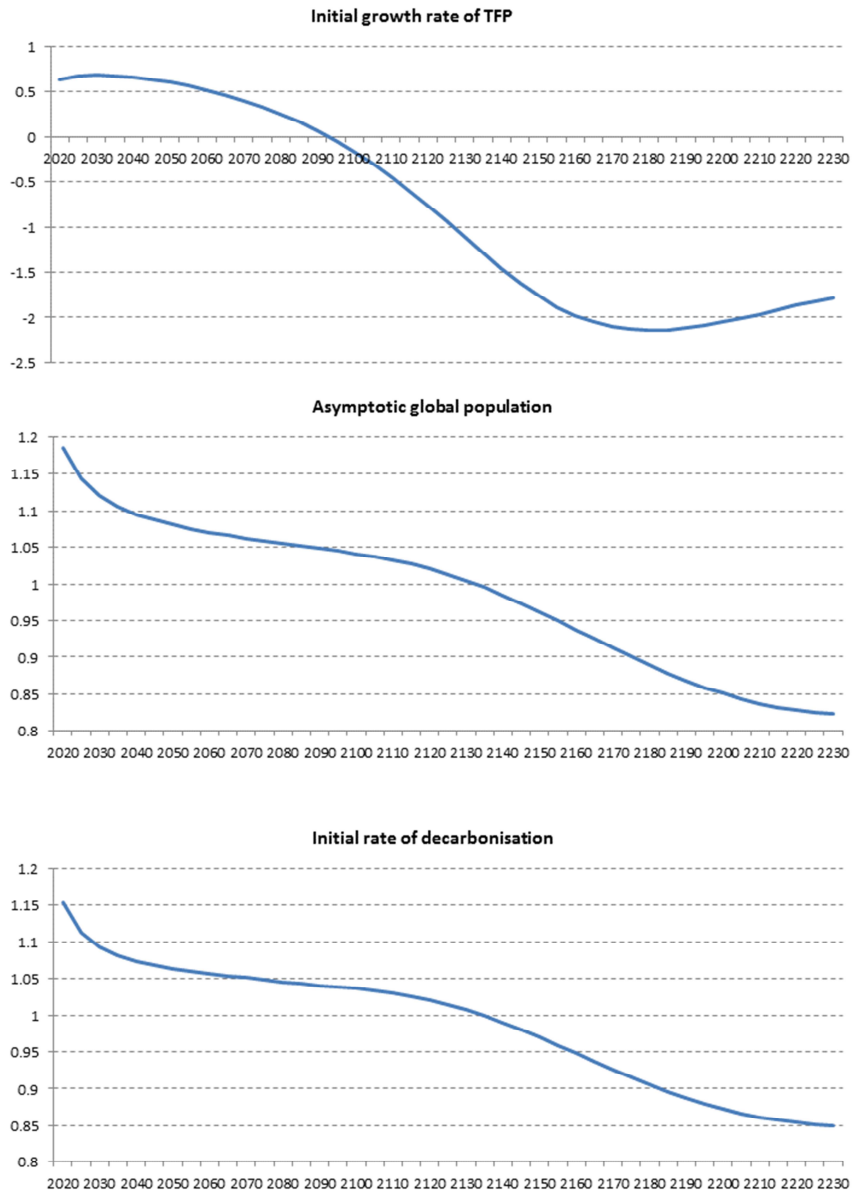

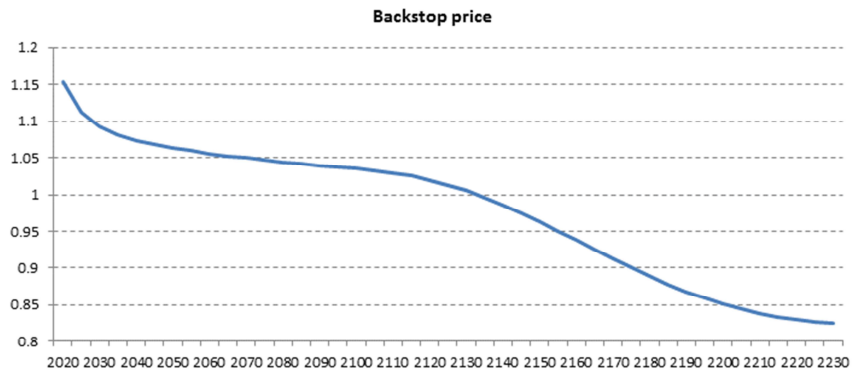
Backstop price
1.2
1.15
1.1
1.05
1
0.95
0.9
0.85
0.8
2020 2030 2040 2050 2060 2070 2080 2090 2100 2110 2120 2130 2140 2150 2160 2170 2180 2190 2200 2210 2220 2230

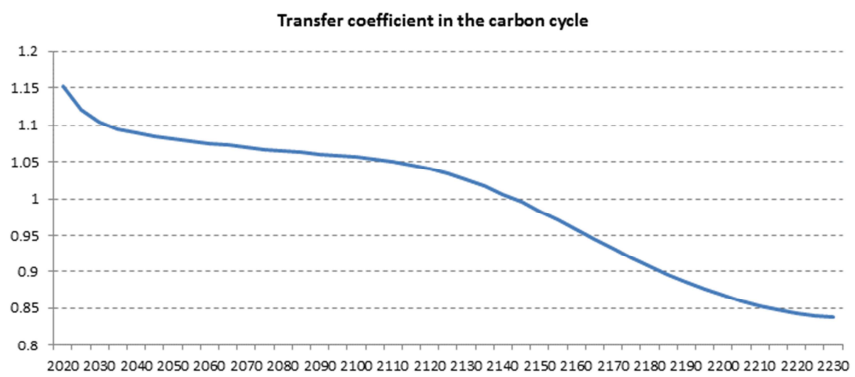
Transfer coefficient in the carbon cycle
1.2
1.15
1.1
1.05
1
0.95
0.9
0.85
0.8
2020 2030 2040 2050 2060 2070 2080 2090 2100 2110 2120 2130 2140 2150 2160 2170 2180 2190 2200 2210 2220 2230

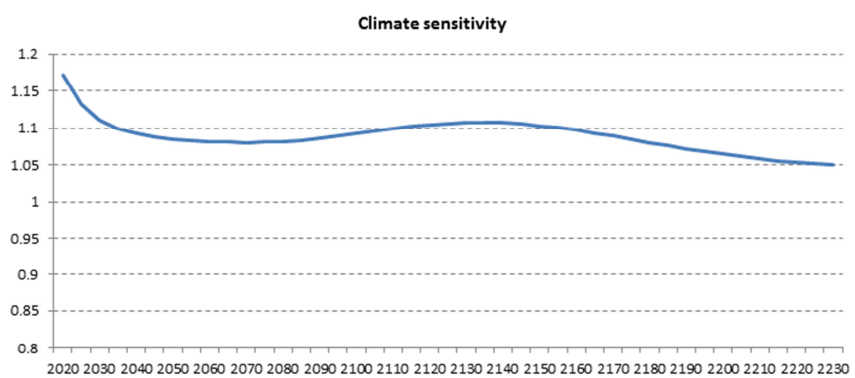
Climate sensitivity
1.2
1.15
1.1
1.05
1
0.95
0.9
0.85
0.8
2020 2030 2040 2050 2060 2070 2080 2090 2100 2110 2120 2130 2140 2150 2160 2170 2180 2190 2200 2210 2220 2230

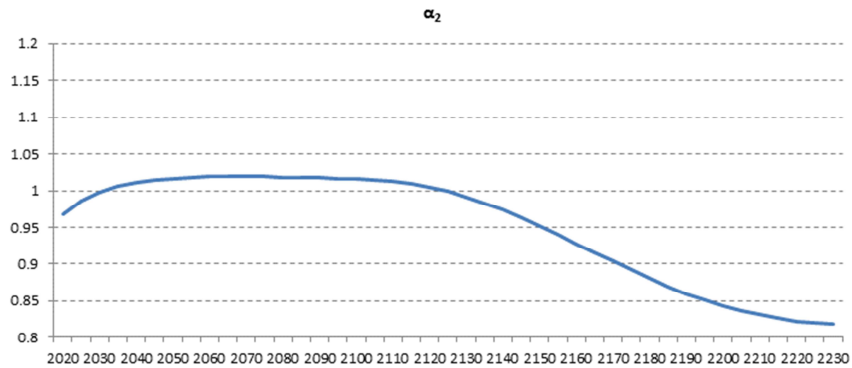
α_2
1.2
1.15
1.1
1.05
1
0.95
0.9
0.85
0.8
2020 2030 2040 2050 2060 2070 2080 2090 2100 2110 2120 2130 2140 2150 2160 2170 2180 2190 2200 2210 2220 2230

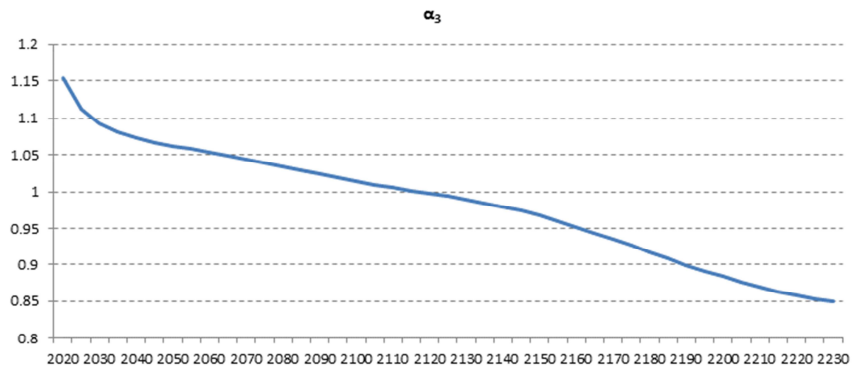
α_3
1.2
1.15
1.1
1.05
1
0.95
0.9
0.85
0.8
2020 2030 2040 2050 2060 2070 2080 2090 2100 2110 2120 2130 2140 2150 2160 2170 2180 2190 2200 2210 2220 2230

6 Nota Bene: a large climate beta is in fact good news for proponents of strong climate mitigation

The climate β is positive and large. This means that when estimating the social cost of carbon, one should use a rate larger than the risk free rate to discount the flow of future expected marginal damages. This effect unambiguously reduces the social cost of carbon. This is because fighting climate change has no hedging/insurance value for risk borne by future generations. On the contrary, it increases the risk they will face. On the face of it this is bad news for proponents of strong and immediate action to reduce greenhouse gas emissions.

But beware, this is not the end of the story! Remember that the climate β is the elasticity of climate damages with respect to changes in aggregate consumption: $\mathbb{E}\left[B_t \left| c_t\right.\right] = c_t^{\beta_t}$. Thus, in the absence of any uncertainty about future consumption, a large β is linked to a large benefit in a growing economy. Moreover, since the benefit is convex in $\ln c_t$ when β_t is positive, the uncertainty affecting future log consumption raises the expected benefit. The two effects work together to raise the expected future benefit to be discounted. More precisely, given the assumptions set out in Section 2, we have that the unconditional expectation of the future benefit equals

$$\mathbb{E}B_t = c_0^{\beta_t}\mathbb{E}e^{\beta_t x_t} = c_0^{\beta_t} e^{(\beta_t\mu + 0.5\beta_t^2\sigma^2)t}. \tag{24}$$

With constant β, $\mathbb{E}B_t$ is exponentially increasing in t when the trend of growth μ is positive. Moreover, the larger is β_t, the larger is the growth rate of the expected benefit. The intuition is as follows. The elasticity of benefits with respect to changes in consumption has two reinforcing effects on $\mathbb{E}B_t$. First, if trend growth is rapid, highly elastic investments will benefit more from economic growth. Second, the benefit is a convex function of the growth rate x_t of consumption. By Jensen's inequality, the uncertainty affecting economic growth raises the expected benefit. Because this convexity is increasing in the elasticity β_t, this effect is increasing in β_t. The combination of these two effects may dominate the discounting effect. Indeed, combining equations (5) and (24) implies that

$$\text{NPV} \quad = \quad \sum_{t=0} c_0^{\beta_t} \exp\left[(-r_f + \beta_t\left(\mu - \gamma\sigma^2\right) + 0.5\beta_t^2\sigma^2)t\right].$$

This is increasing in β_t if β_t is larger than $\gamma - (\mu/\sigma^2)$. This result is summarized in the following proposition.

Proposition 3. *Consider an asset with maturity-specific constant betas, i.e., an asset whose future benefit $B_t|_{t\geq 0}$ is related to future aggregate consumption $c_t|_{t\geq 0}$ in such a way that for all t there exists $\beta_t \in \mathbb{R}$ such that $\mathbb{E}\left[B_t\,|c_t\right] = c_t^{\beta_t}$. Under the standard assumptions of the CCAPM, the value of this asset is locally increasing in β_t if it is larger than the difference between relative risk aversion and the ratio of the mean by the variance of the growth rate of consumption.*

Over the last century in the United States, we observed $\mu \approx 2\%$ and $\sigma \approx 4\%$. If we take $\gamma = 2$, as suggested by Kolstad et al. (2014) on normative grounds, this implies that $\gamma - (\mu/\sigma^2) \approx -10.5$. Alternatively, to acknowledge the equity premium puzzle we might take $\gamma = 10$, so that we obtain $\gamma - (\mu/\sigma^2) \approx -2.5$. Because most actions yield benefits with an elasticity with respect to a change in aggregate consumption larger than either of these two numbers, we conclude that the NPV of most investment projects is increasing in their CCAPM β. 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 β on the expected benefit is much larger than its effect on the discount rate, thereby generating a positive effect on NPV.

7 Conclusion

In this paper we have studied the sign and size of the climate β, using both a simple analytical model and an empirically grounded Monte Carlo simulation of the DICE model. Using the DICE model also enabled us to take into account the effects on the climate β of investment. As long as the structure of climate damages is multiplicative, our results strongly suggest that the climate β is positive. In particular, our numerical modelling with DICE suggests it 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 the damage intensity of warming provide a countervailing effect that tends to reduce β, but it is dwarfed by the effect of TFP uncertainty.

Naturally the validity of our numerical estimates is affected by the well-known weaknesses shared by all IAMs (e.g. Pindyck, 2013; Stern, 2013). And

our estimates 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). 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.

If the nature of climate damages is better represented by an additive structure, then our analytical model shows that the conditions required for a positive climate β are stricter. This raises the question of whether climate damages are better represented by an additive or multiplicative structure? The basic assumption embodied in a multiplicative damage structure is of course that damages are a constant fraction of output, for given warming and damage intensity. By contrast, with an additive structure the share of damages in output, for given warming and damage intensity, decreases as output increases, and *vice versa*. Therefore it is related to the so-called 'Schelling conjecture' that developing countries "best defense against climate change may be their own continued development" (Schelling, 1992, p6). The empirical study closest to answering this question about the structure of damages is Anthoff and Tol (2012), which uses the FUND IAM to estimate the income elasticity of damages on the temporal dimension, disaggregated by region and impact type (other IAMs cannot be used for this purpose of course, because they *assume* a multiplicative structure). They found income elasticities ranging from less than minus one to more than one, however in most regions at most times it is greater than zero and often greater than one. There is therefore little support for an additive damage structure, except at very low income levels, where overall damages are dominated by health impacts that fall with development. More research is clearly required on this issue

Understanding the implications of our findings for climate-change economics requires understanding the dual role played by β in determining the NPV of mitigation, as set out in Sections 2 and 6. It is most straightforward to observe that positive β implies the future benefits 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 $r = r_f + \beta\pi$. The first approach consists in using the systematic risk premium π that has been observed in markets, for instance in the United States over the last century, which has been around 5% (see Gollier (2012), chapter 12). For a project with approximately a unit β, 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, $\sigma^2 = 0.3$ percents on average over the period 2015-2230, so we obtain a risk premium of only 0.6 percents 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 β 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 4), 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, 2007b). Although including these dimensions in our model is beyond the reach of this paper – a new concept of β will need to be developed to accommodate these features – we are inclined to accept this position. We then conclude that a large positive climate β 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, a large positive β implies at the same time a large expected benefit, and a high rate at which to discount it, with an ambiguous overall effect. However, we have shown in Section 6 in the Gaussian framework that NPV is *increasing* in β in its plausible domain, whether one estimates it with a market- or model-based approach. Thus, the large climate β estimated in this paper justifies using a social cost of carbon that is greater than when discounting expected marginal damages at the risk-free rate. It is wrong to think that a negative climate β is the friend of strong action to reduce carbon emissions.

References

Aalbers, R. (2009): "Discounting investments in mitigation and adaptation: a dynamic stochastic general equilibrium approach of climate change," Tech. rep., CPB Netherlands Bureau for Economic Policy Analysis.

Allen, M. R., D. J. Frame, C. Huntingford, C. D. Jones, J. A. Lowe, M. Meinshausen, and N. Meinshausen (2009): "Warming caused by cumulative carbon emissions towards the trillionth tonne," *Nature*, 458, 1163–1166.

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.

Anthoff, D. and R. S. J. Tol (2012): "Schelling's conjecture on climate and development: a test," in *Climate Change and Common Sense: Essays in Honour of Tom Schelling*, ed. by R. W. Hahn and A. Ulph, Oxford: Oxford University Press, 260–274.

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. Quï¿œrï¿œ, 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.

Daniel, K., R. Litterman, and G. Wagner (2015): "Applying asset pricing theory to calibrate the price of climate risk," .

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

Goodwin, P., R. G. Williams, and A. Ridgwell (2015): "Sensitivity of climate to cumulative carbon emissions due to compensation of ocean heat and carbon uptake," *Nature Geoscience*, 8, 29–34.

Howarth, R. B. (2003): "Discounting and uncertainty in climate change policy analysis," *Land Economics*, 79, 369–381.

——— (2009): "Discounting, uncertainty, and revealed time preference," *Land Economics*, 85, 24–40.

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.

——— (2014): "Summary for Policymakers," in *Climate Change 2014: Impacts, Adaptation, and Vulnerability. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, 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. Cariï¿œo 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. Schlï¿œmer, 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, H. D., N. P. Gillett, P. A. Stott, and K. Zickfeld (2009): "The proportionality of global warming to cumulative carbon emissions," *Nature*, 459, 829–832.

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 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.

Schelling, T. (1992): "Some economics of global warming," *American Economic Review*, 82, 1–14.

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.

Van den Bijgaart, I., R. Gerlagh, L. Korsten, and M. Liski (2013): "A simple formula for the social cost of carbon," .

Weitzman, M. (2007a): "A review of the Stern Review on the economics of climate change," *Journal of Economic Literature*, 45, 703–724.

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

Weitzman, M. L. (2007b): "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.

Zickfeld, K., M. Eby, H. D. Matthews, and A. J. Weaver (2009): "Setting cumulative emissions targets to reduce the risk of dangerous climate change," *Proceedings of the National Academy of Sciences*, 106, 16129–16134.