

De Bruin, Kelly C.; Krishnamurthy, Chandra Kiran

Working Paper

Optimal climate policy with fat-tailed uncertainty: What the models can tell us

ESRI Working Paper, No. 697

Provided in Cooperation with:

The Economic and Social Research Institute (ESRI), Dublin

Suggested Citation: De Bruin, Kelly C.; Krishnamurthy, Chandra Kiran (2021) : Optimal climate policy with fat-tailed uncertainty: What the models can tell us, ESRI Working Paper, No. 697, The Economic and Social Research Institute (ESRI), Dublin

This Version is available at:

<https://hdl.handle.net/10419/237968>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Working Paper No. 697

March 2021

Optimal climate policy with fat-tailed uncertainty What the models can tell us

Kelly de Bruin^{a,b} and Chandra Kiran Krishnamurthy^{c,d}

We present a modification of the most commonly used integrated assessment model (IAM) of climate change (DICE-2016), AD-DICE2016, which is designed to address three key aspects of climate-economy models: treatment of uncertainty, the use of more appropriate utility functions, and including adaptation policies to climate change. These modifications ensure that two of the key difficulties identified with IAMs, the choice of the risk aversion parameter and the underestimation of damages, are also directly addressed. The use of a bounded (Burr) utility function ensures that the model is able to appropriately assess the effects of parameters whose distributions have “fat tails”. Uncertainty is accommodated via the state-contingent approach enabling us to include more state (seven) and control variables (four) than recursive derivatives of DICE. Our approach to uncertainty ensures that the optimal climate policies account for outcomes in every possible state, unlike the Monte Carlo approach. Our treatment of uncertainty is extensive: eight parameters are allowed to be random, with distributions –many “fat tailed”– identified using current knowledge. Our model suggests that uncertainty regarding damages and climate sensitivity are key drivers of climate policy. We also find that uncertainty leads to increases in both optimal mitigation and adaptation, with adaptation and mitigation reacting differently to uncertainty over different parameters. Finally, our estimates of the social cost of carbon are larger when uncertainty is allowed for and significantly affected by adaptation.

* Keywords: Climate Change, Uncertainty, Integrated Assessment, Risk Aversion, DICE

* JEL Classification codes: Q54, Q00, D61, D90, C61, C63

Corresponding author: kelly.de.bruin@econ.umu.se, kelly.debruin@esri.ie

a: Economic and Social Research Institute (ESRI), Dublin, Ireland

b: Trinity College Dublin, Dublin, Ireland

c: Department of Forest Economics, Swedish University of Agricultural Sciences, Umeå, Sweden.

d: Center for Environmental and Resource Economics (CERE), Umeå

1 Introduction

Climate change remains the pre-eminent global environmental policy challenge. To enable climate policy setting by appropriately balancing the needs of economic growth with minimising climate-related damages to society, a framework integrating the essence of economics and climate science is essential. Integrated Assessment Models (IAMs), which were developed to model the causal chain of climate change as completely as possible, provide this framework and have been extensively employed by governments and international organisations for evaluating climate change policies.

The relevance of IAMs for climate policy making has been questioned recently (see e.g. [Pindyck \[2013,0\]](#), [Heal \[2017\]](#)), due largely to the following four lacunae: First, the modeling of damages, and in particular, the damage function, are widely acknowledge as being lacking (see e.g. [Pindyck \[2013\]](#), [Hanemann \[2009\]](#), [Stern \[2013\]](#), [Dietz and Asheim \[2012\]](#)) in key dimensions, including estimation techniques (which often rely on outdated and even extrapolated data, and only include a limited number a climate impacts, see e.g. [Tol \[2008\]](#), [Burke et al. \[2015\]](#), [Tol \[2018\]](#)). Furthermore, impacts have been aggregated (across sectors of the economy and regions of the world) due both to lack of data and for model simplicity and computational tractability. Moreover, impacts in IAMs are generally modelled as GDP impacts (so called “level effects”), whereas modelling impacts in terms of GDP growth rate, damage to stocks (capital, labour), and effects on TFP or utility would arguably be a more accurate reflection of reality ([Fankhauser \[2005\]](#), [Wing and Lanzi \[2014\]](#)). Second, the option of adaptation is often either not included at all or only simplistically accommodated, ignoring regional differences in adaptation ability ([de Bruin and Dellink \[2011\]](#), [Tol \[2008\]](#)). Third, IAMs often use unbounded utility functions, and the most commonly used functional form, the Constant Relative Risk Aversion (or CRRA) turns out to be unsuited for the task at hand (see e.g. [Weitzman \[2009\]](#), [Ikefuji et al. \[2011\]](#), [Pindyck \[2013\]](#), [Ikefuji et al. \[2020\]](#)). Fourth, and finally, uncertainty pervades every component of the complex chain of cause-and-effect linking economic output to climate impacts. The effects of uncertainty are difficult to both understand from first principles ([Lemoine \[2021\]](#)) and to appropriately account in the IAMs ([Morgan and Dowlatabadi \[1996\]](#), [Peterson \[2006\]](#), [van Asselt and Rotmans \[2002\]](#), [Weitzman \[2009\]](#), [Pindyck \[2013\]](#), [Stern \[2013\]](#), [Ikefuji et al. \[2020\]](#)).

Given the relevance of integrated frameworks such as the IAMs for climate policy making, many alternative approaches to addressing the issues highlighted above have been explored in the literature¹: a few studies have explored the beneficial effects of more appropriate utility functions ([Ikefuji et al. \[2013\]](#), [Millner \[2013\]](#), [Ikefuji et al. \[2020\]](#)); many studies explore different ways of accommodating uncertainty, be it parametric (using the rather common “Monte Carlo” approach), recursive ([Traeger \[2014\]](#), [Crost and Traeger \[2013\]](#), [Cai et al. \[2015\]](#), [Ikefuji et al. \[2020\]](#)) or a mix of both ([Pizer \[1999\]](#)); a few studies explore whether, and to what extent, the combination of uncertainty (including the effect of so-called “fat

¹ The literature involving advances on many of these fronts are too large to easily review. Nonetheless, we mention a few specific studies here; [Glanemann et al. \[2020\]](#) implement damages based on recent econometric estimates by [Burke et al. \[2015\]](#) and [Piontek et al. \[2019\]](#) implement climate impacts through different channels. [de Bruin et al. \[2009a\]](#), [Bahn et al. \[2019\]](#), [Bosello et al. \[2010\]](#) include adaptation as a policy option.

tailed” distributions of key parameters) and welfare functions and frameworks affect climate mitigation policy (Ackerman et al. [2010], Dietz and Asheim [2012], Ikefuji et al. [2020]). However, very few studies integrate many of these questions, and fewer still consider the question of adaptation to climate change in these frameworks.

The main purpose of this paper is to develop an extension of the most widely used IAM, DICE-2016R2 (detailed in Nordhaus [2018] and often simply called “DICE 2016”), with a view to addressing some of the key challenges detailed above. To this end, the model we develop, AD-DICE2016, extends DICE-2016R2 in the following ways: (i) explicitly allowing for adaptation to climate change, a key component of any climate policy package, in addition to mitigation; (ii) accommodating uncertainty over a large number of parameters (eight), with distributions for each chosen based upon best available data and estimates, and “fat-tailed” distributions allowed in key parameters; (iii) illustrating the use of a more appropriate bounded utility function, which not only allows for more sensible consumption evaluation but also for more appropriate risk profiles at low consumption values. Using the model so-developed, we shed light on three key aspects of climate policy: (1) What effect does uncertainty in parameters exert on optimal climate policies? (2) How do adaptation and mitigation relate to each other, and how does uncertainty affect them both? and (3) How is the Social Cost of Carbon (SCC) affected by uncertainty? While many of the questions we address have been studied—in isolation—in the past, our contribution here lies in analyzing these questions in a richer and unified framework.

DICE-2016 and many recent models derived recently from it (e.g. Traeger [2014], Crost and Traeger [2013]) have two control variables, mitigation and capital investments. AD-DICE2016, which is built on the deterministic version developed in de Bruin et al. [2009a], de Bruin [2011], allows for explicit adaptation to climate change,² leading to an increase in the number of policy (control) variables to four, namely capital investments, mitigation, reactive (flow) adaptation and proactive (stock) adaptation. Proactive adaptation involves long term investments, which occur in anticipation of climate change, whereas reactive adaptation, evidently, occurs in reaction to realized climate change. Thus, in contrast to current approaches to modifying DICE to accommodate uncertainty that rely on simplifying certain dimensions of DICE (Traeger [2014], Crost and Traeger [2013], Cai et al. [2015]), typically by reducing state space, our approach is to *enrich* the basic DICE, leading both to enlarged state space (from six to seven states) and an increase in the number of controls (two to four). In consequence, our model is also able to highlight the relative roles of adaptation and mitigation under uncertainty, itself an interesting topic of research and of much policy relevance. Apart from one study using a framework with one uncertain parameter (Felgenhauer and de Bruin [2009]), adaptation and its interaction with mitigation under uncertainty has been virtually unexplored.³

² Adaptation to climate change refers to social and economic changes which limit the amount of damage associated with a certain level of climate change (Smit et al. [2001]).

³ Prior analytical work (e.g. Ingham et al. [2007]) finds that uncertainty leads to increased adaptation and reduced mitigation, which is an issue our model can shed light on.

Uncertainty in IAMs, which arises from a variety of sources, is challenging to effectively accommodate, leading to a lack of clarity regarding its implications for policy. Integrating uncertainty into IAMs represents a major challenge and has generated a large number of studies with many perspectives and approaches. The two main approaches in the literature differ in terms of whether uncertainty is resolved “all at once” or whether it evolves and persists over time. The first approach, termed “ex-ante” in [Crost and Traeger \[2013\]](#), is often associated with the “Monte Carlo” approach. In this approach (which is widely used in applied studies), uncertainty is not intrinsically a part of the decision making process, allowing one to “draw” the presumptive unknown parameters from their known (or estimated) distributions, and averaging the resultant optimal paths from many different draws. When used in conjunction with commonly applied normal distributions for unknown parameters, this has often led to the conclusion that uncertainty does not substantially affect optimal climate policies (see e.g. [Hope \[2006\]](#), [Richels et al. \[2004\]](#), [Anthoff et al. \[2009\]](#), [Nordhaus \[2008\]](#), [Anthoff and Tol \[2013\]](#), [Hof et al. \[2008\]](#), [Roughgarden and Schneider \[1999\]](#)).

Given the incoherent optimal policies that result when applying this method, the recent literature appears to have largely abandoned the search for optimal policies and used this method solely to investigate a range of outcomes under uncertainty for an exogenously given policy, often focusing on the resulting uncertainty in the social cost of carbon (SCC) as a metric of policy stringency ([Dietz and Asheim \[2012\]](#), [Gerlagh and Liski \[2016\]](#), [Ackerman and Stanton \[2012\]](#), [Hof et al. \[2008\]](#), [Anthoff and Tol \[2013\]](#), [Nordhaus \[2018\]](#)). A novel approach within the Monte Carlo literature is explored in [Gillingham et al. \[2018\]](#), which applies 6 IAMs (DICE, FUND, GCAM, MERGE, IGSM, and WITCH) in an extensive Monte Carlo analysis. This work examines parametric as well as model uncertainty and is also novel in the sense that it estimates surface-response functions to parametric uncertainty.

The second approach (called the “ex-post” approach) consists of modeling uncertainty as being “persistent”, being an intrinsic part of the decision making process, and each period’s decision must be made prior to realization of uncertain aspects of the model. A variety of interesting and new concepts arise here, including learning and the value of information, which have recently been explored in the literature (e.g. [Gerlagh and Liski \[2016\]](#), [Karp and Zhang \[2006\]](#), [Lemoine \[2021\]](#)). This way of introducing uncertainty often leads to recursive solution approaches, typically the use of stochastic dynamic programming (SDP), which provides a framework for accommodating different sources of uncertainty and can also help investigate important aspects not highlighted here (e.g. non-exponential discounting). A major challenge with this approach, however, relates to computational feasibility with regard to the dimensions of the state space encountered in IAMs (dimensions larger than four are computationally very challenging). These challenges have recently been dealt with either by state-space reduction of DICE (e.g. [Traeger \[2014\]](#), [Cai et al. \[2015\]](#)), enabling efficient computation, or by a regression-based simulation approach (the LSMC) enabling more rapid computations ([Ikefuji et al. \[2020\]](#)) with a few uncertain parameters.

In essence, current methodological choices to integrating uncertainty in optimal decisions making in IAMs have been to either apply the intrinsically incoherent Monte Carlo approach or to reduce state space dimension of IAMs rendering the recursive SDP solution numerically feasible (or to use a rela-

tively involved LSMC setting when only a few parameters are uncertain). We use an approach that may be considered a complement to the SDP approach. This approach is coherent from a decision making perspective while being more computationally tractable and straight forward than the SDP approach, and feasible with larger state spaces than the DICE2016R. We use a modified ex-ante approach to uncertainty, termed “state-contingent” in [Pizer \[1999\]](#), with uncertainty arising from parametric uncertainty (i.e. key model parameters are uncertain with known distributions). In our setting, for each of (say) 1000 “states of the world”, possible values of uncertain parameters are drawn from respective (known or estimated) distributions, with each state having a specified probability. The policy maker then maximizes the objective function, which is a weighted (by state probability) sum of utility in each state, and chooses his control variables. Thus, the optimization is *ex-post*, given the known range of uncertain outcomes (and can be seen as a mild form of preference aggregation across states).

[Pizer \[1999\]](#) shows that this approach to uncertainty is more natural and makes a substantive difference to policies compared to the more common ex-ante (Monte Carlo) approach. Since the policy maker optimizes taking into consideration different states of the world, computational feasibility is an important issue and is directly related to the size of the state and control space. Computational feasibility is maintained, however, both due to the smaller state space—the model developed there is not a typical IAM—as well as the use of a closed form approximately optimal consumption rule (implying in particular that the recursive optimization part is never carried out)⁴. Our approach differs from both the Monte Carlo, with optimization carried out ex-post, and the SDP, with a larger state space (seven instead of a maximum of four) and greater number of controls (four instead of two). We also outline a way of accommodating preference heterogeneity (differences in risk preferences) in a manner consistent with existing dynamic welfare frameworks.

We allow eight parameters identified as relevant within an IAM framework to be uncertain. These parameters pertain to five core aspects of the model: economic uncertainty, climatic uncertainty, damage uncertainty, population uncertainty and uncertainty over risk preferences. Also departing from much of current literature (which focuses largely on normal distributions), we allow for a range of plausible distributions for parameters, including those with skew and/or heavier tails. Previous analyses allowing for heavy-tailed distributions focused entirely on uncertainty surrounding either climate sensitivity and damages ([Ackerman et al. \[2010\]](#), [Dietz \[2011\]](#), [Ackerman and Stanton \[2012\]](#), [Anthoff and Tol \[2014\]](#))⁵ or damages, emissions-output ratio and technological efficiency ([Ikefuji et al. \[2020\]](#)). Our study is oriented towards developing an IAM that is capable of addressing many of the drawbacks of treatment of uncertainty in IAMs and is broader in scope than prior studies, covering larger sets of parameters (eight in total) while simultaneously choosing an optimal policy that takes into account parametric uncertainty.

⁴ We note that the approach in [Pizer \[1999\]](#) may be considered a hybrid, in that there is also persistent or structural uncertainty in the model (via a productivity shock)

⁵ In the FUND model applied in [Anthoff and Tol \[2013\]](#) all parameters are varied and hence the fat tails found in the Monte Carlo analyses in FUND are a result of combinations of parameter draws, rather than an assumed fat tailed distribution as input.

In most IAMs, including DICE (Nordhaus [2018]), WITCH (Bosetti et al. [2006]) and MERGE (Manne et al. [1995]), welfare is represented by CRRA utility that, in combination with heavy-tailed distributions for consumption, can imply infinite expected utility and infinite marginal expected utility, which is the conceptual crux of Weitzman [2009]’s “Dismal Theorem” applied to essentially deterministic IAMs. While convenient for many reasons, this form of the utility function is unsuited for the application to climate change, as has been repeatedly emphasized (see discussion in section 2.2). In this context, Ikefuji et al. [2011,0] (and Millner [2013], for a more analytical model) indicate that once the convenience of the CRRA form of utility is discarded, and an appropriately bounded utility function is used, heavy tailed distributions leading to low consumption can be accommodated. We adopt the so-called Burr or Pareto utility function, which at typically observed consumption levels behaves like the CRRA function whereas at more extreme (low) utility levels, behaves like the exponential function (see Ikefuji et al. [2013] and Ikefuji et al. [2020]), instead of the more common CRRA form.⁶

Our results suggest that uncertainty exerts a significant effect on optimal climate policy: mitigation and adaptation are both sensitive to uncertainty, with uncertainty regarding the damage function and climate sensitivity exerting the greatest effect. Broadly speaking, climate policies become more aggressive with uncertainty, meaning that they are shifted forward in time: the SCC rises more steeply over time with uncertainty and full mitigation occurs earlier. Furthermore, mitigation and adaptation are more sensitive to uncertainty in different parameters, and mitigation in particular is very sensitive to uncertainty. Monte Carlo approaches to accommodating uncertainty not only underestimate the effects of parameter uncertainty but are also “incoherent”, and suggest lower mitigation and adaptation levels than the no-uncertainty case.

The paper is structured as follows: Section 2 presents the details of the model, including a description of the utility and welfare functions, and the necessity of a dynamic welfare framework. Section 3 describes the distributions used for key parameters and compares them with those used in the recent literature. Section 4 presents the key results of the model and Section 5 concludes with a discussion on the implications of our findings in the broader context of climate change policy.

2 The AD-DICE2016 model

Here we present a very brief description of the key aspects of AD-DICE2016. AD-DICE2016 replicates the DICE2016-R2 model in terms of the description of economic production and growth, the climate system, emissions and mitigation. We refer the reader to the Supplementary Appendix for complete model equations. Here we focus on the elements of AD-DICE2016 that differ from DICE-2016R2 and

⁶ We note that the only alternative function that has been used is by Dietz and Asheim [2012]. That study however uses a non-utilitarian welfare measure, rendering comparisons of pure functional forms in a discounted utility framework—such as ours—rather difficult. Alternatives to the power function have usually been the log utility function (e.g. Gerlagh and Liski [2016]), which imposes a stronger form of CRRA, fixing risk aversion (inter-temporal elasticity of substitution) at unity, and recursive utility (or Epstein-Zin preferences, Traeger [2014], Crost and Traeger [2013]), which (at least as used thus far) has similar features to the CRRA.

discuss the inclusion of adaptation, the utility function and welfare representation, and finally our method of implementing uncertainty.

2.1 Adaptation

Adaptation is an important policy option that is often overlooked in the IAM literature. Only in the last decade or so have IAMs started to include adaptation as a policy option (e.g. AD-DICE (de Bruin et al. [2009a]), AD-RICE (de Bruin [2011]), AD-MERGE (Bahn et al. [2019]) and AD-WITCH (Bosello et al. [2010])). However, adaptation has thus far never been explored with uncertain parameters in an IAM setting.

The AD-DICE2016 model extends the original DICE-2016 model by incorporating both reactive adaptation (modeled as a flow variable) and proactive adaptation (modeled as a stock variable). The costs and benefits of flow adaptation are realized within the same period, whereas stock adaptation requires a build up of adaptation stock which reduces damages in the future, creating a stream of benefits. Examples of reactive adaptation measures are crop changes, the use of mosquito nets and air conditioning. Examples of proactive adaptation are R&D into new crop species, building of sea walls, and the expansion of irrigation infrastructure. By incorporating both mitigation and adaptation as policy choices, the AD-DICE2016 model creates the opportunity to study both mitigation and adaptation strategies simultaneously. A detailed description of the AD-DICE/AD-RICE modelling of adaptation is provided in de Bruin [2011].

AD-DICE2016 is calibrated based on the assumption of optimal adaptation in the DICE2016 model. Hence the net damages in DICE-2016 represent a combination of residual damages (damages after adaptation) and adaptation costs for the optimal level of adaptation. In this sense, AD-DICE2016 with adaptation applied optimally replicates the DICE-2016 baseline results and AD-DICE2016 with both adaptation and mitigation applied optimally replicates the DICE-2016 optimal (mitigation) results. The AD-DICE2016 model disaggregates the net damage function of DICE-2016 into residual damages and adaptation costs, and the relevant functional forms are calibrated based on the literature detailing the costs and benefits of adaptation. A gross damage function is estimated first, representing climate change damages before adaptation measures are taken, where damages (expressed as a fraction of gross output) are defined as a function of temperature change from 1900 (TE_t):

$$GD_t = \alpha_1 T_{AT,t} + \alpha_2 T_{AT,t}^{\alpha_3},$$

where $\alpha_2 > 0$ and $\alpha_3 > 1$. Residual damages, i.e. damages after adaptation, depend on both the gross damages (GD_t) and the aggregated total level of both adaptation types (PT_t) as follows:

$$RD_t = \frac{GD_t}{1 + PT_t}.$$

Both forms of adaptation are imperfect substitutes for each other, and are aggregated using a Constant Elasticity of Substitution (CES) function. Together with the residual damages function, this describes how adaptation expenditures reduce the damages caused by climate change:

$$PT_t = \gamma_1 \cdot (\beta \cdot SAD_t^\rho + (1 - \beta) \cdot FAD_t^\rho)^{\frac{\gamma_2}{\rho}},$$

where SAD_t is the total amount of adaptation capital stock at time t and FAD_t is the amount spent on reactive adaptation in period t . Furthermore, $\rho = (\sigma - 1)/\sigma$ is the elasticity of substitution between stock and flow adaptation, with $\beta \in [0, 1]$, $\rho \in (-\infty, 1]$ and $\gamma_1, \gamma_2 > 0$. Adaptation capital stock is accumulated in the same manner as conventional capital stock:

$$SAD_{t+1} = (1 - \delta_k)SAD_t + IAD_t,$$

where δ_k is the depreciation rate and IAD_t are the investments in stock adaptation. The total adaptation costs in each period are thus

$$PC_t = FAD_t + IAD_t.$$

The model is calibrated so that in the optimal run the total adaptation costs plus residual damages equal the net damages of the DICE-2016 model.

2.2 Adaptation and mitigation interactions

Adaptation and mitigation represent different approaches to reducing the impact of climate change on the economy. An understanding of how these policies differ in terms of when and how and they affect climate impacts will be helpful in interpreting our results. We explore these differences briefly here. We recall that a key difference between adaptation and mitigation is that adaptation can reduce damages associated with a given level of climate change without limiting climate change itself, whereas mitigation limits climate change (and helps avoid future damages). When higher uncertainty regarding climate impacts are likely at higher levels of climate change, mitigation, which can prevent large temperature increases, has a higher potential to limit these large uncertainties. In consequence, these two policy choices also differ in the timing of cost outlays and benefits reaped. For mitigation, there is a long delay between the initial cost outlays and the resulting reduction in climate impacts.⁷ In contrast, benefits from adaptation are realised in a shorter time frame: Flow adaptation benefits are felt immediately (in the same period these costs are incurred) and stock adaptation benefits are reaped from the period after investments in adaptation capital are made onward.

To illustrate these differences explicitly, we implement an equal increase in spending on each alternative at a point in time and examine the flow of benefits resulting from this one-time increase in outlay. More specifically, we first compute the unrestricted optimal policy (optimal mitigation and flow/stock

⁷ Mitigation reduces current emissions, reducing current concentration of GHG emission in the atmosphere, in turn lowering the increase in atmospheric temperature over time, thus finally reduces climate change impact.

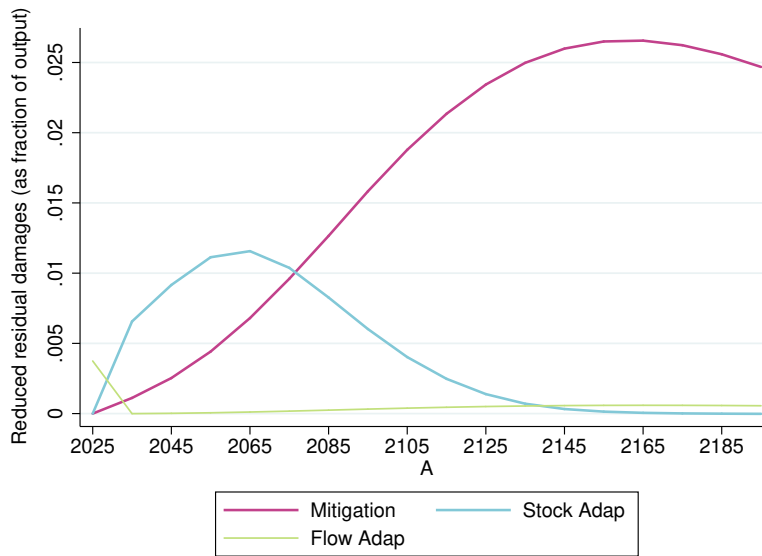


Figure 1: Benefits in terms of reduces damages over time for a one shot increase in mitigation, flow adaptation and stock adaptation spending of 7% of GDP in the period 2025-2035.

adaptation levels and costs), and then implement an identical and exogenous increase in spending on each alternative for one time period only (2025-2035, a ten year period that represents one time-step of the model). In all cases the other policy options are fixed at their optimal levels already computed. In each case, the change in the residual damages over the model horizon that results from this change illustrate when (and for how long) the benefits of this spending are realized. The increase considered (to 7% of GDP) is in all cases larger than the optimal.

The resulting reduction in residual damages (in percentage of GDP) for each policy are displayed in fig. 1. As evident from the figure, mitigation benefits increase initially and peak around 2150, after which the reduction in damages remains constant for the next few decades. Thus, while mitigation benefits accrue from the period after the expenditures occur, the bulk of the benefits are realised in generations to come. In the case of flow adaptation, benefits are only felt in the period in which they are enforced. For stock adaptation the realised benefits are highest in the period after investment and reduce sharply after several periods, but continue up to 2100.⁸

Adaptation and mitigation policies may be considered both complements to, and substitutes for, one another. They are complements in the sense that their different characteristics ensure that both are likely to form non-trivial part of any optimal climate policy, since adaptation may be used to more effectively address challenges in the shorter term while mitigation has the advantage of providing a stream of benefits into the future. They may also be considered substitutes in the sense that they compete for the same resources and an increase in mitigation is likely to reduce the marginal benefits of adaptation (and vice

⁸ The difference in the time profile of benefits make the choice between adaptation and mitigation dependent on the discount rate chosen. Clearly, higher discount rates favour adaptation over mitigation and vice-versa, as has been illustrated in the literature see e.g. de Bruin et al. [2009b] and de Bruin et al. [2009a].

versa). These mechanisms have been illustrated in previous literature [de Bruin et al. \[2009b\]](#) and [de Bruin et al. \[2009a\]](#).

2.3 Utility Function and its Parameterisation

As briefly discussed in section 1, we use the Burr utility function in AD-DICE2016. The challenges of using unbounded utility functions when assessing problems with long time horizons (such as climate change) and possibly unbounded outcome distributions have been highlighted in e.g. [Weitzman \[2009\]](#) and [Pindyck \[2013\]](#). Given that many studies have highlighted the benefits of moving away from the restrictive CRRA-type utility for climate change in particular, we highlight the benefits of using a more appropriate (in a specific sense, see below) utility function here. The specific form of the utility function used is

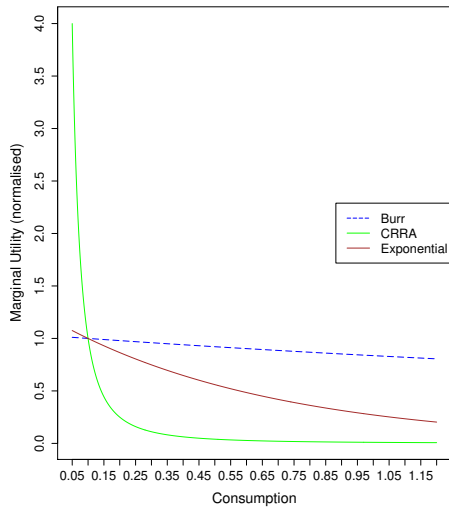
$$U_b(c) = 1 - \left(\frac{\lambda}{c + \lambda} \right)^\kappa, \quad \lambda, \kappa > 0, \quad (1)$$

in contrast to the CRRA-type utility function used in DICE:

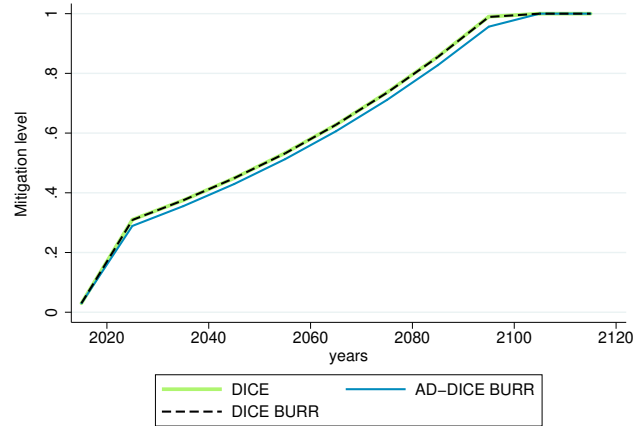
$$U_p(c) = \frac{c^{1-\mu} - 1}{1-\mu}, \quad \mu \neq 0, 1. \quad (2)$$

Where λ and κ are the Burr utility parameters and μ is the elasticity of marginal utility. The Burr utility function in eq. (1), detailed in [Ikefuji et al. \[2013\]](#) (where it is called “pareto utility”), is a two-parameter utility function with several interesting properties of which we briefly mention the following; it is bounded from above and below (by 0 and 1 respectively), and in general behaves like the CRRA utility function remote from the origin and like the exponential utility function near the origin (as is evident from fig. B.2 for consumption close to zero). As a result, it is better behaved than the commonly used CRRA utility function when consumption is close to zero, and is more appropriate when considering risk with heavy tails.⁹ For AD-DICE2016, we calibrate the Burr utility function so as to resemble the behaviour of the CRRA utility function used in DICE-2016R2 at typically observed consumption levels. We choose the parameterization of $k = 0.602$ and $\lambda = 7.593$ for the Burr function to obtain very similar utility function to that in DICE-2016R2 (with $\mu = 1.45$) over almost the entire range of consumption.

⁹ Unlike the CRRA utility function, it also has (i) bounded relative risk aversion coefficient, as can be verified from the expression $RRA(c) = \frac{c(\kappa+1)}{c+\lambda}$, implying $RRA(0) = 0$ and $RRA(\infty) = \kappa + 1$; and (ii) bounded marginal utility, as can be easily seen from $U' = \frac{\kappa\lambda^\kappa}{(c+\lambda)^{\kappa+1}}$, implying $U'(0) = \frac{\kappa}{\lambda}$ and $U'(\infty) = 0$. Indeed, it also has bounded absolute risk aversion coefficient, see [Ikefuji et al. \[2013\]](#) for details.



(a) Marginal utility



(b) Optimal mitigation (emissions reduced as fraction of total emissions)

Figure 2: Comparing Burr and CRRA utility functions

As fig. 2a shows, marginal utility¹⁰—which plays a key role in choice of optimal climate policy—differs significantly between the CRRA and the Burr, exponential utility functions: the Burr reduces slowly from its maximum of 1 (similar to the exponential); the CRRA begins at very large values (being unbounded) close to zero consumption and falls to very low values for consumption levels at which the other two still have relatively large values of MU.

Figure 2b shows the optimal mitigation in DICE-2016R2 with CRRA utility, DICE2016R with our specified Burr Utility function as well as the AD-DICE2016 with the Burr utility function.¹¹ It is evident from the figure that the difference between the CRRA and our parameterization of the Burr utility function is negligible. Furthermore, the difference between the AD-DICE2016 mitigation path (with optimal adaptation and Burr utility) does not differ significantly from that for the original DICE-2016R2 (with Burr utility).

We note that the choice of a bounded Burr utility function enables us to allow for, and explore, the possibility of consumption falling to very low values when distributions with heavier tails than the normal are used. Over the “normal” range of model outcomes, the Burr and the CRRA yield very similar quantitative and qualitative outcomes, as also indicated in Ikefuji et al. [2011] who use a simplified and

¹⁰ The figure presents normalised values of marginal utility, where each utility function’s marginal utility is normalised by its corresponding value at $c = 0.1$. This normalisation merely ensures that very large values of CRRA marginal utility are comparable to those of the other two utility functions (whose marginal utility, recall, are bounded). It mechanically follows from the normalisation that the marginal utility of all three utility functions are identical at $c = 0.1$.

¹¹ Different versions of DICE have been used as a benchmark in different studies in the literature. Our interest is in enriching DICE in a variety of ways, and consequently, we benchmark our model and results with, where possible, those of the latest version of DICE, DICE-2016R2. The differences between this model and its predecessors is detailed in Nordhaus [2017].

two-period version of DICE with the CRRA and Burr utility and in [Ikefuji et al. \[2020\]](#), who use the DICE2016R with a recursive approach to uncertainty (with what they term “light tails”).

2.4 Welfare Framework

We use a utilitarian approach to social welfare, similar to the analysis in [Pizer \[1999\]](#); however, our approach will be somewhat different from the rather static approach used there. To see why an explicit framework for social welfare is necessary, we begin by observing that there are essentially two types of uncertainties involved in our framework: the first is induced by uncertainty over (non-preference) parameters; the second is over preferences, induced by uncertain preference parameters. If the only uncertainty were over non-preference parameters, then no aggregation and social welfare function is needed, one can simply work with population-weighted expected utility (as in many stochastic variants of the DICE model). The introduction of uncertainty over preference parameters necessitates some form of aggregation over different “types” of preferences in different “states of nature” (in preference terms).

The simplest form of preference aggregation is provided by Harsanyi’s theorem on Social Aggregation for a static setting, and is used in [Pizer \[1999\]](#) (see section A for a brief overview). In view of that static setting being unsuitable for a dynamic framework, we use instead newly developed aggregation principles for dynamic decision-making. In the context of dynamic choice under uncertainty, aggregation of preferences is a rather tricky affair. We only point out a few issues involved and the implications for our setting, and refer to [Zuber \[2011\]](#), [Jackson and Yariv \[2015\]](#) for details regarding recent developments. We note that our framework closely follows that in [Zuber \[2011\]](#).

Society is composed of M “individuals” (they will be defined formally below) and uncertainty over consumption arises from parametric uncertainty. An uncertain consumption stream arising from a generic distribution (or ‘lottery’) L is written as $\{c_t^L\}_{t=0}^\infty$.¹² Individuals compare an infinite stream of uncertain consumption yielded by L using expected utility, represented using the usual VonNeuman-Morgenstern (vNM) utility function,

$$U_i(L) = \mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t u_i(c_t) \right], \quad (3)$$

with u_i and δ the individual i ’s Bernoulli utility function and discount factor respectively.

This is the usual expected infinite stream of utility, with the usual properties. The important question to be answered is: what are the assumptions needed on social ordering of the “individual” consumption streams (in eq. (3)) to yield an inter-temporal version of Harsanyi’s aggregation theorem for the static case (presented in equation A.1). [Zuber \[2011\]](#) (Proposition 3) shows that if social aggregation respects a Paretian principle and if the social planner’s preferences are also vNM, then social preferences satisfying

¹² We abstract away from many issues, such as that related to the cardinality of the consumption set, and direct the reader to [Zuber \[2011, §2\]](#) for a fuller account of the theory. Note also that individual i ’s consumption stream being possibly different from individual j ’s, consumption streams are to be written $\{c_{t,j}^L\}_{t=0}^\infty$. For notational simplicity, we do not index the consumption streams by individual indices and also omit the superscript L . Note also that all expectations are implicitly assumed to be w.r.t the lottery L , written as \mathbb{E}_L .

reasonable properties (stationarity and history-independence) are of the weighted utilitarian form¹³

$$\bar{W}(L) = \mathbb{E}_L \left[\sum_{t=0}^{\infty} \delta^t W(c_t) \right], \quad (4)$$

with $W(c_t) = \sum_{i=1}^M u_i(c_t)$, the “social Bernoulli utility function”.¹⁴ In other words, society evaluates an uncertain consumption stream using an expected vNM-Bergson type of SWF,¹⁵ with period social welfare an unweighted sum of individual utilities in that period. Note in particular that the entire framework (in (3) and (4)) requires that all individuals and society evaluate utilities at the same discount rate, $\rho = \frac{1-\delta}{\delta}$.¹⁶ In summary, we allow preference parameters other than the discount rate to vary and fix the discount rate.

Preference uncertainty is accommodated by allowing the utility functions of “individuals” to differ in a specific way, whose essence is easily explained: if M is the number of unique combinations of the preference parameters, then each preference parameter combination represents an “individual”, and eq. (4) represents a coherent way of deriving an explicit welfare function that accounts for preference uncertainty. We focus on the coefficient of relative risk aversion as the preference parameter of interest, and allow the functional parameter of Burr utility that controls the RRA to have a “fat tailed” distribution (the Pareto, see section 3.4).

2.5 Accommodating parameter uncertainty: The State-contingent Approach

Choosing the optimal policy under uncertainty means that decisions in each period must be made prior to realization of the (random) parameters of the model. The two prominent reasons for uncertainty arising in climate change applications relate to state evolution (e.g. random shocks to state variables) and parameter uncertainty (true value of key parameters being unknown), where parameters can take values from known distributions. Stochastic dynamic programming (SDP), is the approach that can best capture stochastic behavior when uncertainty arises from recursively expressed state transition. This approach, which suffers

¹³ In the dynamic version of welfare from eq. (4), note that the weights, ω_i in equation in A.1 are all set to 1.

¹⁴ This equation can be rewritten as follows to aid understanding:

$$\bar{W}(L) = \sum_{t=0}^{\infty} \delta^t \mathbb{E}_L [W(c_t)] = \sum_{t=0}^{\infty} \delta^t \sum_{i=1}^M \mathbb{E}_L (u_i(c_t)).$$

Thus, the evaluation over time may, to a first approximation, be thought of as a sum of the static evaluation. Of course, two differences are apparent: first, that the weights are now unity, for each individual and second, each individual can differ in attitude to risk and to consumption over time (since the u_i vary by individual i) but not over the discount factor, δ .

¹⁵ Note that, strictly speaking, this SWF is actually *ordinal* since no conditions rendering utilities comparable (as in the static setting) are imposed. This is evident since utilities are unweighted.

¹⁶ This is a very strong requirement and cannot be dispensed with; Jackson and Yariv [2015] show, in a setting with no uncertainty, that not imposing this requirement leads to plans which are present-biased and history-dependent in a particular way, while Zuber [2011] cites examples to show that this represents inter-temporal indifference to consumption inequality.

from the so-called “curse of dimensionality”,¹⁷ has been applied with modified forms of IAMs in e.g. Traeger [2014], Cai et al. [2012] and Cai et al. [2015].

The approach we use is a form of *non-recursive* stochastic programming, suited for cases where the distribution of outcomes (in our case parameters) is independent of both the actions/controls (mitigation and adaptation, in our case) and state variables. Thus, the key conceptual difference between the recursive formulation (via SDP) and the state-contingent approach is that the states of the world, which are known in advance (i.e. can be enumerated), cannot be affected by actions taken (e.g. policies chosen) in previous periods. As a result, this is not a framework suited for application to a model where the state variable is uncertain (e.g. is subject to random evolution), since in such an instance the evolution of the state variable must depend upon its current value and upon current policy. On the other hand, it is ideally suited for a framework, such as ours, where parameters are uncertain, since the decision maker’s actions (e.g. mitigation policies) cannot affect the (denumerable) parameter distribution (e.g. distribution of climate sensitivity).¹⁸

In the state-contingent approach, the objective function is a weighted sum of the utility in each state; in other words, utility for each possible realization of the parameters is computed for that period, and the weighted sum of this utility is the maximand. An advantage of this approach over SDP is that stochastic models with larger state space can be accommodated; on the other hand, the drawbacks include an inability to accommodate recursive features of uncertainty (as already referred to) or other interesting features such as (structural) learning regarding parameters. Despite its advantages, to our knowledge only one study, Pizer [1999], has applied this approach to integrated assessment models. Pizer [1999] presents a framework for determining optimal climate change policy under uncertainty and his model is essentially a modified DICE, using contingent utility (subject to the proviso in footnote 3).

Implementation of this approach is conceptually straight forward: the “uncertain states of the world” are the possible values of the parameters. Assuming that the parameters are independent, a draw (of size J) is obtained from their marginal distributions in time period 1.¹⁹ It is not known however which state will actually occur, so the policy maker has to take the multiple potential futures into account. Optimisation is ex-post, in that the optimal policy takes into account the many possible states of the world. Thousands of states are used to simultaneously determine the optimal policies for each scenario. A large number of states of the world is necessary to explore the realistic policy consequences of uncertainty, else important aspects related to interaction between the parameter space and the model may be overlooked. A simplified outline of the utility structure is given below (the complete version is provided in the Supplementary

¹⁷ The “curse” refers to the fact that the computational time and space required for solving dynamic programming problems numerically rises exponentially with the dimension of the state space (in terms of the number of state variables). In practice, computing optimal policies becomes impractical with more than four state variables.

¹⁸ In the literature on stochastic IAMs, we are unaware of any study that deviates from the assumption that parameters are fundamentally a fixed number independent of our actions; it is the actual value of the parameter that is unknown. In models with learning, this true-but-unknown value is potentially knowable after some experimentation.

¹⁹ This is essentially the “naïve” random sampling of individual parameters. We do not pursue more involved sampling approaches (e.g. the Latin Hypercube) since it is not clear that in cases such as ours, with as many as 8 variables random, and with only a single draw of them all, the LH approach provides any benefits.

Appendix), which includes two illustrative state variables, one for capital stock (K_t) and one for climate-related components ($M_{AT,t}$):

$$\begin{aligned} & \max_{I_t, MU_t, IAD_t, FAD_t} \frac{1}{N} \sum_{s=1}^N W_s \\ & s.t. K_{t+1} = f(K_t, I_t) \\ & M_{AT,t+1,s} = f(M_{AT,t,s}, M_{UP,t,s}, I_t, MU_t, \theta) \end{aligned} \tag{5}$$

where K_t is capital stock, $M_{AT,t,s}, M_{UP,t,s}$ represent the the atmospheric and upper ocean layer concentration of CO_2 at time t and state s , W_s is welfare, index $s = \{1, \dots, N\}$ represents states of the world (at most N), $\theta \in \{\theta_1, \dots, \theta_N\}$ is a draw of an uncertain parameter vector from the set of all possible values while investment I_t , mitigation MU_t , flow adaptation FAD_t and stock adaptation IAD_t are the decision variables. Welfare in state s is defined as the discounted sum of utility $u()$

$$W_s = \sum_{t=1}^T R^t [u()],$$

where $u(.)$ is the Burr utility function defined in section 2.3 and R^t is the discount factor. We note that our model has seven state variables, with the four controls mentioned here.

3 Quantifying parameter uncertainty

Our model includes over 70 parameters. Based upon a variety of factors, including availability of knowledge enabling parameterisation, we choose eight parameters to treat as uncertain.²⁰ These parameters are listed in table 3, along with typical estimates of location and scale values and distributions used in the literature. We categorise these parameters into five groups, based on their role in the model: Economic (denoted ECO), climatic (CLIM), damages (DAM), population (POP) and risk (RISK). Each group of parameters corresponds to different components of the climate-economy cycle. As already alluded to above, most modeling studies incorporating parameter uncertainty typically consider only a subset of the parameters below as being uncertain at a time (Ackerman et al. [2010], Dietz [2011], Dietz and Asheim [2012], Ackerman and Stanton [2012], Roughgarden and Schneider [1999], Anthoff et al. [2009], Ikefuji et al. [2020]). In essence, that approach can be considered a moderate ‘‘robustness check’’ of the model rather than as an exploration of the implications of joint parameter spaces on climate policy that we undertake.²¹

²⁰ We have chosen to not include uncertainty over the costs associated with adaptation and mitigation, chiefly since the data for these are largely unknown, making the task of assigning probability distributions more challenging.

²¹ The recent study of Dietz et al. [2017] considers as many (or more) parameters uncertain as we do. However, the focus there is very different, on estimating the (CAPM theory-derived) ‘‘climate beta’’. Consequently, there is no notion of ‘‘optimality’’ in the sense commonly used, much less a discussion of mitigation and policy in the standard IAM sense.

Other studies include a more substantial set of uncertain parameters, but assume normally distributed parameters, around a rather narrow variance (see e.g. Hof et al. [2008], Hope [2006]). Some studies include uncertainty over all model parameters with the drawback of not only applying normal distribution but also assuming equal distributions for all parameters (Anthoff and Tol [2013], Anderson et al. [2014]). Turning to studies most directly related to our work, many of the chosen distributions in Pizer [1999] have compact support while the distributions in Nordhaus [2008,0,0] and Gillingham et al. [2018] are generally normal. This is rather a strong assumption, ruling out “extreme” model outcomes in the expected utility setting underpinning all IAMs. These assumptions have often been criticised, see e.g. Weitzman [2009], Pindyck [2013], Stern [2013], as problematic, being both incoherent in theory and in the resulting dependence of policy choices upon ad-hoc distributional bounds. We carry out a more extensive simulation exercise, differing from the existing literature in using distributions reflecting the underlying nature of the parameters, a few of which have “fat tails”. In any case, prior studies investigating the effects of parameter uncertainty upon climate policy lack a coherent link between uncertainty and climate policy, while studies in which uncertainty is central to climate policy (e.g. Traeger [2014], Cai et al. [2012], Ikefuji et al. [2020] and Cai et al. [2015]) have focused on model structures which are simplified in certain dimensions and focus on only a few parameters and do not undertake an extensive exercise in identifying the effect of different sources of uncertainty.


We note that our choice of distributions and their parameters draws upon the (by now large) literature in economics that has developed around understanding how uncertainty affects different aspects of climate policy. In particular, our choice of distributions follows either from the characteristics required of certain phenomena (e.g. “fat-tailed” distributions for climate sensitivity) or what is considered most sensible, largely for reasons of non-negativity (e.g. truncated normal and triangular). Similarly, the distributional parameters are drawn based upon best estimates in the scientific (for e.g. climate sensitivity, transfer coefficient) or economic (for e.g. the coefficient of risk aversion) literature. In light of the independence of parameters,²² we draw from respective marginal distributions whose choice we turn now to discuss.

3.1 Damage Exponent


The damage function links temperature increases to losses in output and is typically quadratic (e.g. DICE-2016 in Nordhaus [2017] and Nordhaus [2018] and many other integrated assessment models) with the form $\alpha_1 T_t + \alpha_2 T_t^2$. This functional form not only lacks empirical support but is also limiting at higher temperature increases, yielding improbably low damage estimates. For instance, both Ackerman et al. [2010] and Weitzman [2011] have shown that with Nordhaus’ values of the damage exponent, a 5°C warming results in a loss of utility equivalent to 6% of output, which is unreasonably low (see also Stern [2013, §3.2.1]), while Hanemann [2009] provides a detailed argument of why damages in the DICE

²² Given that varying each parameter in a large-dimensional-parameter-space setting is computationally challenging (since the number of runs will be the product of the number of simulations for each parameter), some form of simplification is essential to render computations feasible. We make the (mild) assumption that parameter distributions are independent of one another.

models are under-estimated. Weitzman (Weitzman [2012, §3]) suggests a damage function with the power exponent taking a value of 6.



(a) Triangular distribution for damage exponent in AD-DICE2016 and Ackerman et al. [2010].



(b) Damage function for different values of damage exponent.

Figure 3: Properties of the damage function in AD-DICE2016.

We use instead the functional form $\alpha_1 T_t + \alpha_2 T_t^{\alpha_3}$, which allows for a wider range of damage possibilities, and allow the exponent α_3 to be a random variable. In view of the non-negativity of this parameter, we choose the triangular distribution, which is often preferred under these conditions (Kotz and Van Dorp [2004]), and has also been used in many studies on climate change (e.g. Ackerman et al. [2010]). As to the support of this distribution, a lower bound of 2 is conventional and the upper end of the support can be fixed by understanding the sensitivity of damages to changes in the exponent: we plot damage (fraction of output) as a function of temperature increase for different values of the damage exponent α_3 (fig. 3b). An exponent of 7 leads to an almost vertical damage curve, and we choose a max of 5 (where damages are near-vertical). In this case, gross damages would be approximately 2.5% of GDP (excluding adaptation) at a 2°C temperature increase, 62% of GDP at a 4°C increase and 186% at a 5°C increase. Finally, a choice of center = 3.3 yields a mean of 3.4, this is consistent with the base value in the AD-DICE2016 model²³

²³ We note a few properties of the triangular distribution, a continuous distribution with finite support. Denoting by a and $b < \infty$ the support, and by c the “most likely value” (i.e. mode of the distribution), of a random variable X with the triangular distribution, the density function is (see eq (1.6), Kotz and Van Dorp [2004]):

$$f(x|a,b,c) = \begin{cases} \frac{2}{b-a} \frac{(x-a)}{(c-a)} & \text{if } a \leq x \leq c \\ \frac{2}{b-a} \frac{(c-x)}{(b-c)} & \text{if } c \leq x \leq b \end{cases},$$

with mean $\mu = \frac{a+b+c}{3}$ (see eq (1.22), Kotz and Van Dorp [2004] for an expression for the variance).


Turning to a brief comparison with other studies, Nordhaus [2018] introduces uncertainty over the coefficient on the second-order term (α_2), using a normal distribution with mean 0.00227 and standard deviation of 0.00135. Ackerman et al. [2010] uses the same functional form as we do for the damage function, with the exponent a random variable following the triangular distribution with $\min = 1$, $\max = 5$ and $\text{center} = 2$, yielding a mean of 2.95. Dietz and Asheim [2012] use the damage function $\alpha_1 T_t + \alpha_2 T_t^2 + \alpha_4 T_t^7$, introducing an ad-hoc high-order term to capture the extreme non-linearity at large temperature levels and a normal distribution for its coefficient α_4 (an effect that may be accommodated instead by higher powers of T). Pizer [1999] uses a quadratic damage function and models uncertainty by using a discrete distribution with 5 values for the coefficient of the second-order term. Finally, Ikefuji et al. [2020] use a quadratic damage function (with $\alpha_1 = 0$) with a student-t distribution for α_2 with the d.o.f. parameter of 10 (parameter values for σ are not reported, with μ possibly being unbounded).

To summarize, the form of the damage function we use, along with the range of the distribution for the damage exponent, allows for stronger damages than in many prior studies, and responds to the criticism in Stern [2013] regarding substantial underestimation of damages. The choice of a commonly used damage function, along with a transparent way of allowing damages to scale with temperature change, also addresses, to an extent, the criticisms of IAMs in Pindyck [2013] regarding unrealistic and ad-hoc choice of damage functions.


3.2 Climate Sensitivity

The climate sensitivity parameter, which governs the relationship between atmospheric CO_2 concentrations and the global mean temperature, measures the equilibrium response of temperature to a doubling of atmospheric CO_2 concentration. It is a key parameter because it captures the climate system's response to variations in the Earth's radiation balance caused by changes in atmospheric CO_2 . If climate sensitivity is high, relatively small changes in CO_2 concentrations could produce very significant warming effects, with severe consequences for both human society and planetary ecology.

The IPCC's Fourth Assessment Report has compiled a number of recent estimates of the climate sensitivity parameter, and an inspection of the 18 probability density functions in box 10.2 in Meehl and Stocker [2007] shows that most (if not all) of them have a positive skew and a "fat tail" of high estimates. The IPCC concludes that the climate sensitivity is "likely" (66 – 90% chance) to be in the range of 2 – 4.5°C with a best estimate of about 3°C, and is "very unlikely" (< 10% chance) to be less than 1.5°C. This means that there is a 5 – 17% chance that climate sensitivity is greater than 4.5°C. A more recent average estimate by fourteen leading climate scientists of the probability of climate sensitivity being above 4.5°C is 23% (Zickfield et al. [2010]). This means that a normal distribution of the climate sensitivity parameter may significantly underestimate the probability of very high values.



(a) Climate sensitivity distributions: Nordhaus [2018], Ackerman et al. [2010], Dietz and Asheim [2012] and AD-DICE2016



(b) Truncated-normal distribution for the transfer coefficient: $\mu = 0.24, \sigma = 0.1$

Figure 4: Density functions of Climate Sensitivity and the Transfer Coefficient

We use a lognormal distribution with $\mu = 1.34, \sigma = 0.5$ chosen to match the estimations compiled by IPCC as closely as possible (see 4a). Most other studies also use log normal distributions with comparable, yet slightly less “fat tailed” distributions, with distributions in studies most directly related illustrated in fig. 4a (see also Weitzman [2012, Table 1]): Nordhaus [2018] and Gillingham et al. [2018] model uncertainty by assigning a log-normal distribution (fit to Olson et al. [2012]) with a mean of 3.10 and a standard deviation of 0.84. Ackerman et al. [2010] and Dietz and Asheim [2012] also use a lognormal distribution. Pizer [1999] uses a discrete distribution with 5 values almost symmetrically centered around the value 3 (1.5, 2.2, 3, 3.7, 4.5). Tables 1 and 2 presents cumulative probabilities for different distributions. To summarize, the distribution for climate sensitivity we choose, unlike much of the prior literature, has tails which are “fat” enough to match the latest scientific estimates.

| Model | $\hat{T} =$ | $2^{\circ}C$ | $4^{\circ}C$ | $6^{\circ}C$ | $8^{\circ}C$ |
|------------------------------|-----------------------------------|--------------|--------------|--------------|--------------|
| Nordhaus [2018] (LN) | | 0.92 | 0.12 | 0.003 | 0 |
| AD-DICE2016 (our model) (LN) | | 0.90 | 0.46 | 0.18 | 0.070 |
| Ackerman et al. [2010] (LN) | $\mathbb{P}_{LN}(T \geq \hat{T})$ | 0.73 | 0.24 | 0.07 | 0.022 |
| Dietz and Asheim [2012] (LN) | | 0.85 | 0.23 | 0.038 | 0.0061 |

Table 1: Cumulative probabilities for different temperature increases.

| | $T < 1.5^\circ$ | $2 \leq T \leq 4.5^\circ$ | $T \geq 4.5^\circ$ | Mean |
|-------------------------|-------------------------|---------------------------|-------------------------------|------|
| IPCC | “very unlikely” <10% | “likely” 66-90% | “cannot be excluded” 5-17% | 3 |
| Nordhaus [2018] | 0.5% | 87% | 5% | 3.13 |
| AD-DICE2016 (our model) | 3.1% | 59.8% | 37.1% | 4.32 |
| Ackerman et al. [2010] | 12.3% | 70% | 17.7% | 3.4 |
| Dietz and Asheim [2012] | 15% | 70% | 3.8% | 3.2 |

Table 2: Cumulative probabilities for IPCC temperature ranges.

3.3 Transfer Coefficient in Carbon Cycle

Many integrated assessment models, including DICE2007, AD-DICE2016, Dietz and Asheim [2012] and Ackerman et al. [2010], model the carbon cycle as having three reservoirs of CO_2 , the atmosphere, upper and lower ocean, with carbon flowing in both directions between adjacent layers. The parameters chosen as key to describe this process is the transfer coefficient between atmosphere and upper ocean; i.e. how the upper ocean absorbs CO_2 from atmosphere.

In our model this transfer coefficient is a random variable which follows a truncated normal distribution $[0,1]$ $\mu = 0.24, \sigma = 0.1$ (see 4b). Dietz and Asheim [2012] use a normal distribution with the same parameters as DICE2007 while for Ackerman et al. [2010], the transfer coefficient is not random. Pizer [1999] uses a simplified version of the carbon cycle model with a discrete distribution (with values 0.5, 0.6, 0.65, 0.7, 0.8) for the retention rate for CO_2 emissions. Nordhaus [2018] does not apply uncertainty to a climate transfer coefficient but to the level of carbon in the intermediate reservoir (in GtC), applying a normal distribution.


3.4 Burr Utility Function Parameter

Recall from 2.3 that the Burr utility is a bounded (between 0 and 1) utility function used instead of the CRRA for its suitability for applications with possibly very low consumption values, and that the Burr utility function parameters κ and λ are chosen to closely replicate the control variable results of DICE-2016. The choice of two key aspects of decision making, risk aversion and discount rate, has generated a large literature; indeed, Pindyck [2013] forcefully criticizes IAMs for the ad hoc choices regarding these aspects. The welfare framework outlined in 2.4, however, allows an interesting way of side-stepping some of the criticism involved in the choice of a coefficient of risk aversion. In particular, we are interested in allowing for uncertainty over risk aversion, with an interpretation of variation in risk preference being indicative of different individual types (as elaborated in 2.4).

We recall that the coefficient of relative risk aversion (RRA) for Burr utility, $RRA(c) = \frac{c(\kappa+1)}{c+\lambda}$ (see footnote 9 for details), depends upon consumption level, unlike for the commonly used CRRA utility. An inspection of the effects of κ (fig. 5a) and λ (fig. B.1b) suggests (as is also evident from an inspection

of the expression for the RRA) that variation in κ induces larger changes in *RRA* than variation in λ ; as a result, we introduce uncertainty over the utility parameter κ . We choose the pareto distribution with parameters $k = 0.5$, $a = 5$, resulting in a mean value of 0.626 and a standard deviation of 0.16.²⁴ The choice of the pareto distribution ensures that relatively large values for κ , implying greater aversion to risk, are not ruled out; on the other hand, the mean and variance parameters chosen ensure that very large values of κ , in particular larger than 3, have a very small probability (of 0.00044), as indicated in fig. 5b. The value of $\kappa = 3$ is chosen since it is evident (from fig. B.1a) that utility for low values of c are indistinguishable for $\kappa > 3$.

For all practical purposes, the *RRA* varies between 1 and 3 (fig. 5a), for values of c even moderately beyond 0. These bounds on the *RRA* is also consistent with the commonly used values (usually for the CRRA utility), which vary from 1 to 3 or 4.²⁵



(a) RRA for varying values of κ ($\lambda = 7.5$)

(b) Density function for κ : Pareto($k = 0.5, a = 5$).

Figure 5: RRA of Burr utility function.

Ackerman et al. [2010] and Nordhaus [2008] use the CRRA utility function with fixed parameters, RRA $\mu = 2$, and (the pure rate of time preference) $\rho = 0.015$, while in the DICE-2016R2, the RRA is reduced to 1.45 (see Nordhaus [2017] for more details). Pizer [1999] also uses a power utility function,

²⁴ We use the standard definition of a pareto distribution: for a random variable X distributed Pareto, the density function is $ak^ax^{-(a+1)}$, with $x \geq k$, $k, a > 0$, k the scale and a the shape parameter (e.g. p.574, Johnson et al. [1995]). The mean, μ of this distribution is computed as $ak(a-1)^{-1}$ for $a > 1$ (see p.577 of Johnson et al. [1995] for computation of variance, σ^2).

²⁵ We note that Ikefuji et al. [2020] use a slightly different parameterisation of Pareto utility for their study: $\kappa = 1.322, \lambda = 0.0108$. Our parameterisation was chosen based upon ensuring that the RRA was reasonably close to the commonly used values for CRRA utility (of 1.5) for consumption ranges far enough away from 0.

but introduces uncertainty in both ρ and the RRA, with a mean value of about 1.1 (and a posterior distribution which appears to be a truncated normal, $TN[0,4]$, see Table 1, §2.4).²⁶

3.5 Other Parameters


We turn next to detailing the distributions of the remaining four parameters. The growth rate of total factor productivity is modelled as decreasing over time, with the initial growth rate defined as a random variable, with a truncated normal distribution on $[0, 1]$, $\mu = 0.076$, $\sigma = 0.033$. Nordhaus [2008], Dietz and Asheim [2012] and Pizer [1999] also use a truncated normal distribution. However, Nordhaus [2018], Gillingham et al. [2018] apply a normal distribution resulting in a mean yearly productivity growth of 1.53% and a standard deviation of 1.12%. Despite the normal distribution being used, the significantly higher s.d. compared to other studies (including this study), will turn out to be important.

The CO_2 -equivalent emission-output ratio declines over time as energy efficiency increases. The initial rate of this decarbonization is a random variable and we use again the truncated normal distribution, with parameters $\mu = 0.0152$, $\sigma = 0.005$. We use a triangular distribution for asymptotic global population, $\min = 9770$, $\max = 13000$, $\text{center} = 10800$, with a mean of 11190 (same as the no uncertainty value), based largely upon the UN report (United Nations (2017)). Dietz and Asheim [2012] and Nordhaus [2008] use a normal distribution while Pizer [1999] and uses a discrete distribution for annual decline of population growth rate. Nordhaus [2018] treats population level as deterministic while Gillingham et al. [2018] uses a normal distribution for population level in 2100 (mean of 12.149 billion and standard deviation of 2.378 billion). For total resources of fossil fuels, which defines the cumulative limits of carbon use, we use a triangular distribution, $\min = 5000$, $\max = 8000$, $\text{center} = 5000$, with a mean of 6000 (same as the no uncertainty value). Nordhaus [2008] and Dietz and Asheim [2012] use a normal distribution instead while it is deterministic in Nordhaus [2018]. An overview of all parameters treated as random in AD-DICE2016 is provided in table 3.


4 Results

We turn next to presenting the results of our numerical simulations. Before doing so, we briefly discuss the details of our simulation: we draw randomly from our parameter distributions 1000 times, creating


²⁶ We note that increased value of coefficient of RRA in the CRRA class of utility functions has two distinct and opposing effects: without uncertainty, μ represents aversion to generational consumption inequality, with larger values implying greater aversion to inequality and acting thereby to *reduce* the value of future consumption (since consumption is growing at a constant rate in DICE); when uncertainty is introduced, μ represents the aversion to risk, and since future consumption is uncertain, larger values of μ tend to increase the value of future consumption. In terms of the social cost of carbon, which is discussed in detail in section 4.5, it is not obvious even for the CRRA functional form which of these two effects dominate (although see Pindyck [2013] for a small analytical exercise). For the functional form we use, and the setting in which uncertainty is introduced, the latter feature is likely to lead to increases in the coefficient of RRA—which is non-constant, recall—to lead to more conservative policies, increasing mitigation and the SCC.




(a) Growth rate of TFP: Truncated Normal, $\mu = 0.076, \sigma = 0.033$



(b) Rate of decarbonisation: Truncated normal, $\mu = 0.015, \sigma = 0.005$



(c) Asymptotic Global Population: Triangular, $a = 9770, b = 10800, c = 13000, \mu = 11500$



(d) Total Resources of Fossil Fuels: Triangular, $a = 5000, b = 8000, c = 5000, \mu = 6000$

Figure 6: Parameter Distributions (density functions shown).

| Description | Group | Unit | “No Uncertainty” Value | Distribution | Parameters |
|--|-------|---------------------------|------------------------|------------------|---|
| Initial growth rate of factor productivity | ECO | Fraction per 10 years | 0.076 | TN[0, 1] | $\mu=0.076, \sigma=0.033$ |
| Rate of decarbonization | ECO | Fraction per 10 years | 0.015 | TN[0, 1] | $\mu=0.015, \sigma=0.005$ |
| Total resources of fossil fuels | ECO | Billions of tons carbon | 6000 | Tri(a, b, c) | $a = 5000, b = 8000, c = 5000, \mu=6000$ |
| Damage exponent | DAM | Number | 3.4 | Tri(a, b, c) | $a = 2, b = 5, c = 3.3, \mu=3.4$ |
| Asymptotic global population | POP | Millions | 8600 | Tri(a, b, c) | $a = 9770, b = 10800, c = 13000, \mu=11500$ |
| Transfer coefficient in carbon cycle | CLIM | Fraction per decade | 0.24 | TN[0, 1] | $\mu=0.24, \sigma=0.1$ |
| Climate sensitivity | CLIM | $^{\circ}C/CO_2$ doubling | 4.32 | LN[1.34, 0.5] | $\mu=1.34, \sigma=0.5$ |
| Burr utility function parameter | RISK | | 0.6 | PAR[0.5, 5] | $\mu=0.6, \sigma=0.16$ |

Notes: 1. Tri(a, b, c) refers to the triangular distribution with range $a, b (< \infty)$ and center c , whose mean μ is computed as $\frac{a+b+c}{3}$ and TN(a, b, μ, σ) refers to the standard form of the truncated normal distribution with truncation points a and b .

2. The “No Uncertainty” column refers to the values of the parameters in the base AD-DICE 2016 model, when all parameters are deterministic.

Table 3: Description of uncertain parameters.

1000 equally likely possible states of the world.²⁷ We then optimise our adaptation (flow expenditures and stock investment), mitigation and capital investments paths over these 1000 states. The damages, production levels and social cost of carbon etc. associated with the chosen path will vary over states. We search for an optimal path using the Conopt algorithm in GAMS, which is a non-linear programming algorithm that uses pivoting based on second order derivatives to find the optimal solution. This approach is rather different from the value iteration approaches to approximate the value function in dynamic programming problems based upon the Bellman equation. One benefit of our approach is that state space reduction is less critical for numerical reasons: recall that our state space dimension (of seven) is larger than the six in DICE-2016R2 and the four in [Traeger \[2014\]](#)). Similar to all non-linear optimization algorithms, a disadvantage is that numerical solutions are increasing in the size of the random draws, meaning that the use of e.g. 10,000 realizations will be numerically practically infeasible.

We focus on five aspects in the discussion of results. First, we present evidence suggesting that, in practice at least, the presence of “fat-tailed” distributions need not raise concerns when applying IAMs for policy analysis. Second, we discuss the effects of different sources of uncertainty. Third, we examine how adaptation and mitigation interact under uncertainty, an aspect not studied thus far in the literature, an investigation of which can yield important insights. Fourth, we evaluate how the popular Monte Carlo approach compares, in terms of coherence of policies, to ours. Finally, we present and discuss the social cost of carbon and how it is affected by uncertainty.

4.1 “Fat-tailed” distributions and climate policy

Weitzman, in his so-called “Dismal Theorem” ([Weitzman \[2009\]](#)) has suggested that policy evaluation (including IAMs) that use unbounded utility functions may be challenged when using “fat-tailed” distributions that can lead to very low consumption, whose marginal valuation must be unbounded. This raises challenges that are both modeling-related (needing ad-hoc lower bounds on the marginal WTP, which determines mitigation) and conceptual (how can such bounds be found and reconciled with a theory of welfare?). This view has been reinforced in particular in [Pindyck \[2013\]](#) and [Pindyck \[2017\]](#), who developed further this line of thinking with particular focus on the Social Cost of Carbon, which we investigate later.

As already mentioned, many studies (e.g. [Nordhaus \[2011\]](#)) suggest that the result is a particular artifact of the use of CRRA functions, whose properties may not be appropriate to use when addressing the challenges posed by climate change. We contribute to this debate by considering whether the choice of utility function will result in different optimal policy pathways under fat-tailed uncertainty. We recall that our use of “fat-tailed” distribution for climate sensitivity, in conjunction with the damage exponent being uncertain (and potentially rather large) means that gross damages can reach up to 99.8% of output. Based purely upon marginal utility considerations, we would expect that the CRRA utility function would

²⁷ We have also run our simulations for 5000 draws and find the results consistent with the 1000-draws presented here. As a 5000 draw scenario takes far longer (about two weeks) to solve, we have chosen to use a 1000 draw across the different scenarios.

result in higher mitigation pathways in order to avoid potential catastrophic damages (entailing collapse of consumption to very low values). As the Burr utility function bounds marginal utilities (and the coefficient of risk aversion), one would anticipate lower optimal mitigation levels.

Figure 7 displays the optimal mitigation and adaptation paths with both the Burr and CRRA utility functions (where both climate sensitivity and the coefficient of RRA have “fat-tailed” distributions). Contrary to our a priori expectations, we find only minimal differences between the optimal mitigation and adaptation paths for these two utility functions. While our simulations used a moderate-sized draw (meaning that a very much larger draw may still lead to difficulties with the CRRA utility), our analysis leads us to suggest that the use of fat-tailed distributions can be accommodated in IAMs and used for policy evaluation, at least when reasonable welfare functions are chosen.

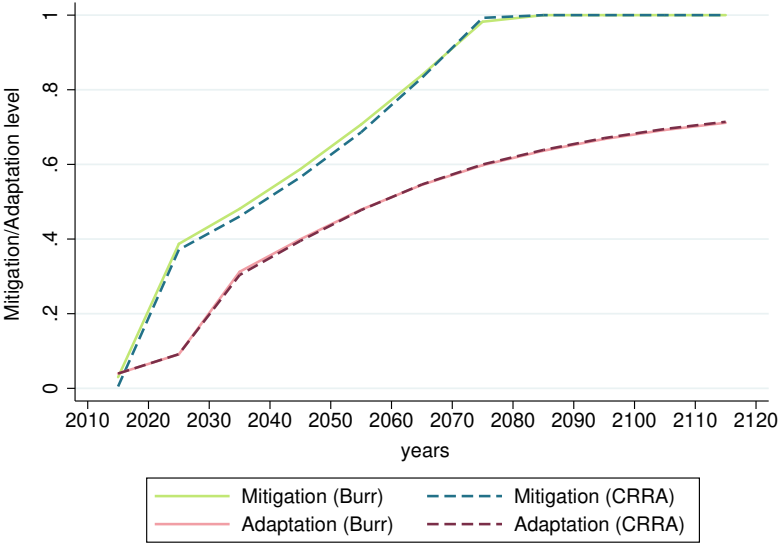


Figure 7: Optimal mitigation and adaptation paths with full uncertainty for Burr and CRRA utility functions) sources of uncertainty.

4.2 Salience of different sources of uncertainty

As detailed in section 3, there are five different sets of uncertain parameters, each corresponding to uncertainty in different components of the climate-economy cycle. In this section, we investigate the degree to which these different sources of uncertainty affect mitigation and adaptation. We also investigate whether adaptation and mitigation are affected differently by these different sources of uncertainty, an investigation that, to our knowledge, has not been carried out before in the literature. Observe that when one set of parameters is set to be uncertain, all others are fixed at their mean values presented in table 3 (except, ob-

viously, for the cases of full and no uncertainty). Similarly, when considering mitigation only, adaptation (stock and flow) is “turned off”, and similar is the case when considering adaptation only.²⁸

The results of our “mitigation-only” model-run are presented in 8a, from which it is evident that different sources of uncertainty affect optimal mitigation paths differently. Following e.g. Nordhaus [2008] and Pizer [1999], we anticipate that uncertainty regarding parameters distributed normally and with linear first order conditions will have no significant effect on the optimal mitigation path. Our results concerning population growth (which satisfies these conditions), confirms this hypothesis, since this does not significantly alter the path of mitigation (path omitted from the figure, as being identical to the no-uncertainty path). When preferences— here the “Burr parameter” κ —are uncertain, the optimal mitigation path is unaffected (path omitted from figure, as being identical to that for the case of “no uncertainty”). Pizer [1999], on the other hand, finds that uncertainty in preference parameters (ρ and RRA) can exert a sizeable effect on optimal policies, with higher levels of mitigation occurring earlier.²⁹ One possible reason for his findings is that ρ is uncertain, which leads to non-linear first order conditions (uncertainty regarding this parameter has an increasing effect on the effective discount rate over time).

Economic uncertainty (total factor productivity growth, decarbonisation rate and available fossil fuels) exerts a sizable effect on economic variables such as output (as in Nordhaus [2018] and Gillingham et al. [2018]) without leading to a discernible effect on optimal mitigation policies. This result suggests that it is important to investigate the effect of uncertainty not only on intermediate variables such as output (as in Nordhaus [2018] and Gillingham et al. [2018]), but also how it affects the optimal policy (which, for our case at least, is minimal).

Turning next to uncertainty regarding climate (climate sensitivity and the transfer coefficient) and damage parameters (the damage coefficient), however, we see a very different pattern. Uncertainty regarding these two key parameter sets leads to sizeable increases in mitigation, with damage uncertainty having a stronger effect than climate uncertainty. In particular, the main effect in both cases is increased mitigation for all time periods, and in particular full mitigation slightly earlier (at 2055) for damage uncertainty than in the case of no uncertainty (at 2065). While the importance of these two parameter sets in driving climate change policies has been emphasized, previous findings on this score have been mixed: more recent studies report that damage uncertainty exerts a sizeable effect on mitigation in a range of models, from the fully SDP approaches (Crost and Traeger [2013]) to modified DICE-based MC approach (Ackerman et al. [2010]) while the same effect was not seen in many of the earlier modeling studies relying on normal distributions (Nordhaus [2008]) or approximate DP approaches even when using (slightly) skewed distributions (Pizer [1999]).

²⁸ We note that mitigation to reduce climate change damages is “switched off”, but mitigation as a means of reducing the use of scarce fossil fuel resources is still available. The level of this mitigation is insignificant except in the case of economic uncertainty, where uncertainty over the amount of fossil fuel resources available results in significantly higher levels of mitigation.

²⁹ Other studies that have considered many parameters uncertain, e.g. Dietz and Asheim [2012], do not provide separate analysis of the importance of uncertainty in different parameter sets.

Similarly, while many studies emphasize the key role of climate sensitivity and damage parameters, particularly with heavy tails, in driving climate policy (Ackerman et al. [2010], Ackerman and Stanton [2012], Dietz [2011]), they do not consider uncertainty over other parameters, making it difficult to identify what aspect drives key insights. In any case, our results support the broad trends in recent findings: that uncertainty in these two parameter sets are of key importance for mitigation. As already mentioned, some recent studies, e.g. Nordhaus [2018] and Gillingham et al. [2018], have emphasised the salience of uncertainty of economic growth, which may partly be explained by the rather high s.d. chosen: in Nordhaus [2018], for instance, this is as large as 74% of the mean, whereas it is only 34% on average for the other four uncertain parameters. Nonetheless, while this does affect economic variables and emissions, damage uncertainty is reported as being the main driver of variation in the social cost of carbon (followed by climate sensitivity and productivity growth). Finally, allowing all parameters to be uncertain (“full uncertainty”) leads to mitigation levels higher than when any particular set of parameters is uncertain, again highlighting the importance of examining multiple uncertainties together.

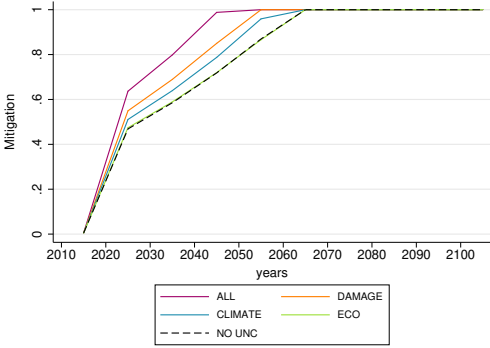
Model results for the adaptation-only case are presented in fig. 8b, which presents impacts post-2035 (differences for periods earlier are relatively small). As for mitigation, uncertainty regarding population and preferences have little effect on optimal adaptation (paths omitted from the figure). It is also evident that uncertainty in damages and all parameters being uncertain (full uncertainty) have similar impacts, signalling that damage uncertainty is a key channel through which uncertainty affects adaptation.

The effects of climate uncertainty on adaptation policies are small overall but exhibit patterns different from mitigation: reductions in adaptation in earlier periods and increases in later periods. This interesting response is driven by two factors: non-linearity in the effect of climate uncertainty (with climate uncertainty leading to large variation in temperature changes with concomitant variations in gross damage across different states-of-world) and the exponentially increasing adaptation cost function (meaning that a unit of increased adaptation costs more than is saved by a reduction in the same unit of adaptation). The latter aspect, in fact, means that when moving from a deterministic model (with an “average” level of gross damage) to an uncertain one with say 1000 states, optimal adaptation may well be reduced. On the other hand, this same shift from a deterministic to a stochastic model, when considering only the effect via gross damages (the former channel identified above) is likely to lead to increased adaptation. In earlier periods of the model runs, the adaptation cost effect dominates the non-linear climate damages effect.

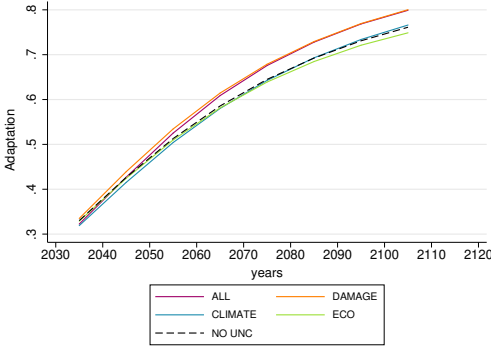
In later periods, the effect on gross damages (which are exponentially increasing) becomes more prominent: carbon-emissions drive temperature increases due to the monotonic increase in economic activity. Due to the considerable time lag between emissions and resulting temperature change (see section 2.2), impacts peak after approximately 50 years, meaning that changes in climate sensitivity exerts a stronger effect in later periods. Furthermore, gross damages increase exponentially in temperature (due to the exponential gross damage function). In later years, the effect of these mechanisms outweigh the effects of the exponential adaptation cost function and the optimal level of adaptation increases.

The economic parameters, total factor productivity (TFP) growth, rate of decarbonisation and available fossil fuels, turn out to have significant effects on adaptation, and we consider the effect of each in turn. When TFP growth is uncertain, average economic growth increases (since TFP growth has exponential effects on production), which in turn has two competing effects: the first is an obvious increase in emissions, warranting greater levels of adaptation, *ceteris paribus*; this is counteracted by the second, involving reduced marginal utility of consumption over time, leading to an increase in Ramsey discounting (consumption discounting), reducing the size of future impacts leading to reduced stock adaptation level. A similar effect operates upon the rate of decarbonisation, i.e. uncertainty leads to an average higher level of decarbonisation, leading to reduced emissions and thus reducing the need for adaptation. Uncertainty of fossil fuel resources will lead to increased mitigation to reduce fossil fuels usage, leading to reduced adaptation (see footnote 28). In summary, economic uncertainty has effects upon adaptation in both directions but overall, we find that the downward pressures dominate, and optimal adaptation level decreases with economic uncertainty.

We also find that, up to the year 2080, adaptation levels are higher when only damages are uncertain than when all parameters are uncertain, largely due to economic uncertainty in the latter case leading to reduced adaptation. These climate and damage uncertainty results are interesting, and illustrate that different sources of uncertainty do not uniformly affect the direction of adaptation levels. These effects are not observed with mitigation, largely since the effects of uncertainty are manifest through different channels.³⁰



(a) Optimal mitigation (without adaptation) with different sources of uncertainty.



(b) Optimal adaptation (without mitigation): different sources of uncertainty.

Figure 8: Adaptation and Mitigation–no interaction.

Mitigation and adaptation are measured as fractions of emission reduction and output respectively.

³⁰ We note that stock and flow adaptation are affected very similarly by parameter uncertainty, and we therefore omit the details (figures are available upon request).

4.3 Adaptation and mitigation interactions under uncertainty

To our knowledge, no previous study using an IAM has explored how adaptation and mitigation policies interact under uncertainty (with the exception of the simplistic analysis in Felgenhauer and de Bruin [2009]). Figures 9a and 9b present optimal levels of mitigation and adaptation under uncertainty in AD-DICE2016. We first examine the effect on mitigation. Comparing mitigation with (fig. 9a) and without (from fig. 8a) adaptation yields several insights: First, mitigation levels are (unsurprisingly) higher without adaptation than with. Second, when adaptation is allowed for, uncertainty over economic parameters does affect mitigation, with optimal mitigation reduced (the ‘ECO’ curve in is slightly below the ‘No UNC’ curve beyond 2065 in fig. 9a, whereas it was identical to the ‘No UNC’ curve in fig. 8a). As previously discussed, uncertainty over economic parameters leads to an increase in both economic growth and de-carbonisation. Absent the possibility of adapting to climate change, mitigation is higher, with full mitigation reached by 2065, while the effects of economic uncertainty only manifest later, post-2065. This is because, as previously discussed, prior to 2065 increased economic growth resulting from uncertainty over economic parameters is balanced by an increase in consumption discounting and decarbonisation, resulting in a negligible effect on mitigation. Post-2065 however, the effects of consumption discounting and decarbonisation dominate that of economic growth, resulting in a reduced need for mitigation. Absent adaptation, however, the need to mitigate post-2065 is nonetheless very high (i.e. marginal benefits of mitigation far outweigh the costs) and mitigation is limited by the fact that the 100% threshold (meaning no carbon storage) is reached. Even when economic uncertainty drives down marginal benefits of mitigation, they still exceed marginal costs, warranting full mitigation.




Figure 9: Optimal adaptation and mitigation for different sources of uncertainty
 Mitigation is in fraction of emissions reduced, adaptation in fraction of damages reduced.

Turning next to adaptation, we compare both adaptation levels and the effect of different sources of uncertainty when mitigation is allowed (fig. 9b) with when it is not (fig. 8b). This comparison leads to the following observations: First, adaptation levels are lower with mitigation than without. This is intuitive, since adaptation must be increased to compensate for increased damages resulting from not mitigating.

Second, a surprising finding is that adaptation levels are higher when damages are uncertain (fig. 9b) than when all parameters are (“full uncertainty”), up to year 2090. This was seen in the case of no mitigation as well, driven by the impacts of economic uncertainty, however, it is amplified when mitigation is introduced. This result arises largely due to the timing of mitigation: with all parameters uncertain, full mitigation (meaning that all emissions are avoided) is reached by 2070, while for the case of uncertainty in damage parameter only, this is reached a decade later. This additional mitigation (between 2070 and 2080) with all parameters uncertain naturally reduces the need for adaptation, lowering the optimal adaptation levels.

The opposite effect is observed for the case of climate uncertainty (fig. 9b), where adaptation levels are below the no uncertainty case before year 2090, and above beyond. This follows from similar lines of reasoning regarding mitigation: here, adaptation is rather low until full mitigation (at 2090, for “CLIMATE”). After 2090, however, non-linearities in the climate system mean that, even with full mitigation, increase in climate damage warrants increased adaptation.

In summary, mitigation and adaptation are affected differently by different sources of uncertainty, and considering both together under uncertainty is important to understanding their optimal provision.³¹

Both mitigation and adaptation are strongly affected by uncertainty in two parameters identified as key in the literature; climate sensitivity and damage exponent. Adaptation is significantly more sensitive to damage uncertainty (and to a lesser extent to uncertainty in economic parameters) than mitigation. Uncertainty over other parameters in isolation, including preference parameters, has no discernible effect on either optimal adaptation or optimal mitigation policies.

4.4 Comparison with the Monte Carlo Approach

The Monte Carlo approach, the most commonly used approach when accounting for uncertainty in IAMs, consists of running the deterministic model with many (here, a thousand) draws of all uncertain parameter. Many previous Monte Carlo analyses (e.g. Nordhaus [2008]) have indicated that uncertainty over various parameters, particularly in combination with the choice of normal distributions for uncertain parameters, do not substantially affect climate change policies. We now compare how adaptation and mitigation under uncertainty from our model (stochastic AD-DICE2016) differs from that using the Monte Carlo approach. Figure 10a compares the optimal mitigation paths under both approaches. The results of the MC approach are presented in the form of the average (over different paths) mitigation level, as is common. The median and the range of Monte Carlo outcomes are also provided as well as the no uncertainty optimal mitigation path. Furthermore the lowest (MC-lower) and highest levels found (MC-upper) in the Monte Carlo runs are also graphed to indicate the range of results. Note that we use the same distributions for uncertain

³¹ Regarding the effect of uncertainty on adaptation costs and investments, we briefly mention key changes (details and figures available upon request). In brief, allowing for mitigation significantly alters these: optimal adaptation costs are half as large, damage-related uncertainty in particular affects adaptation. Overall, sensitivity in adaptation costs and investments to uncertainty is muted when mitigation is allowed.

parameters (and the same number–1000–of draws) for our approach and for the Monte Carlo approach, to keep results comparable.

The key finding of 10a is that the choice of a method of dealing with uncertainty (MC versus our method) matters for climate policy: introduction of uncertainty in parameters is seen to reduce optimal mitigation when the Monte Carlo approach is used, contradicting the findings of the MC approach to DICE in e.g. Nordhaus [2008], where little or no effect is found. The differences, we conjecture, are due both to the choice of parameter distributions as well as the choice of the damage exponent (with non-linear first order conditions) instead of the damage coefficient (with linear first order conditions) as being uncertain. We find that when uncertainty results in a high variation of mitigation outcomes, using the MC approach (and averaging across the resulting mitigation paths) leads to lower mitigation levels. Moreover the mitigation path suggested by the Monte Carlo method and our model also differs: the path suggested by our model is more or less parallel to that of the no uncertainty case while the Monte Carlo optimal path has a very different curvature (and intersects the “No Uncertainty” mitigation curve very early on, starting 2030). This illustrates a problem with Monte Carlo analysis, where the average (or any summary measure) optimal control path is incoherent, from a decision making perspective.

Our findings related to optimal adaptation paths are similar: fig. 10b provides a comparison of optimal adaptation paths (in the absence of mitigation) for our model (AD-DICE2016) and the Monte Carlo approach. It is evident that the adaptation path using AD-DICE2016 is always higher than the no uncertainty case. Using the Monte Carlo method, however, leads to a path lower than the no uncertainty case, similar to the case for mitigation.

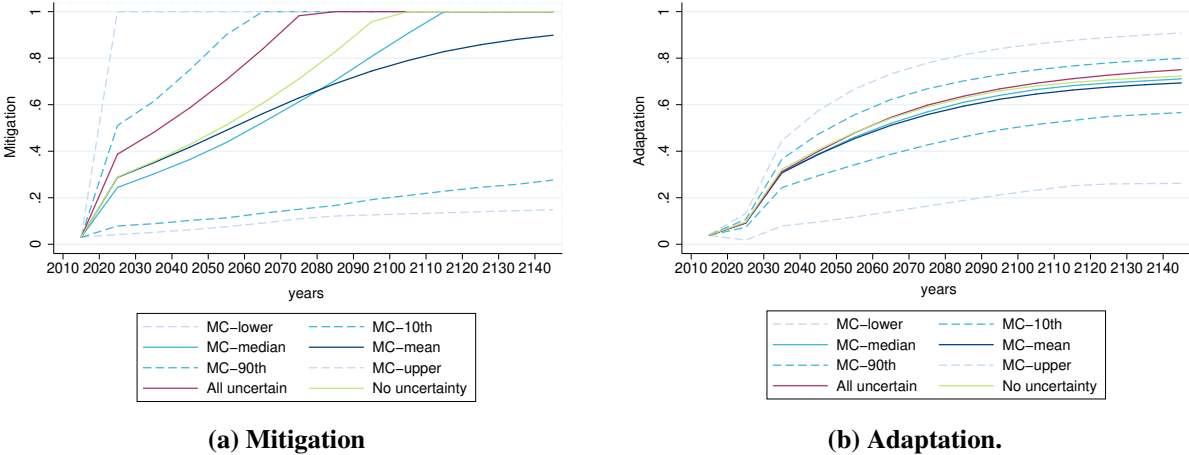


Figure 10: Comparison of adaptation paths and mitigation paths under uncertainty with the Monte Carlo method

Notes: ‘MC’ denotes Monte Carlo method (with 8 parameters uncertain), and “10–th”, “90–th” represent the respective percentiles. “All uncertain”, “No uncertain” correspond to our model (AD-DICE 2016) with all (8) parameters uncertain and no parameters uncertain respectively.

Our results are consistent with those in the study of [Crost and Traeger \[2013\]](#), which provides a comprehensive comparison of the Monte Carlo approach with the approximate DP approach, with only the damage coefficient and damage exponent in turn uncertain. This study reports not only sizeable quantitative differences but also the incoherence of the Monte Carlo approach. In terms of mitigation, it finds that both their ex-post approach and the Monte Carlo (ex-ante) approach result in lower levels of mitigation (when damage coefficient is uncertain). When examining mitigation expenditures, however, this study finds that the Monte Carlo approach suggests *increased* mitigation expenditures, while their ex-post approach suggests decreased expenditures. In summary, the mitigation paths suggested by the Monte Carlo approach need bear no relation to any notion of optimality.

4.5 Climate change costs

Climate change costs are commonly summarised by the social cost of carbon (SCC), and recent discussions related to the effect of policy change using IAMs has often focused on this important component (see e.g. [Pindyck \[2017\]](#), [Weyant \[2017\]](#)). The SCC reflects the shadow price of carbon in terms of climate change damages on different paths (of the control variable). While not the same as an (optimal) carbon tax, it is nonetheless the most important metric of the size of the climate externality.³² The diversity of ways of computing the SCC can make it difficult to compare these estimates across studies. To obviate this issue, we use the approach to computing the SCC following that used in the DICE, which is detailed next for convenience.

The SCC is an estimate of the economic value of the extra (or marginal) impact caused by the emission of one more tonne of carbon at any point in time [Yohe et al. \[2007\]](#). Following [Nordhaus \[2017\]](#), we compute the SCC as

$$SCC(t) \equiv \frac{-\frac{\partial U}{\partial E(t)}}{\frac{\partial U}{\partial C(t)}}. \quad (6)$$

We note that the marginal impact of emissions at time t on utility is divided by the marginal utility value of a unit of consumption in period t , which provides the monetary value of the impacts of an additional unit of emissions. The SCC is estimated within the AD-DICE2016 model and represents a discrete approximation of eq. (6). As in the DICE-2016R2, the SCC is expressed in terms of 2010 US \$ (PPP) per metric ton of CO_2 .

Recall that our model optimizes the controls over the many possible states of the world; as a result, there is one estimate of the SCC for each possible state realised with uncertainty. Following other studies with a similar characteristic ([Crost and Traeger \[2013\]](#), [Traeger \[2014\]](#), [Nordhaus \[2018,0\]](#)), the SCC (US \$ per ton of CO_2) is presented as summaries over all possible states, in particular, the mean and the 10th

³² Indeed, the U.S. government constituted Inter-agency working group on the social cost of carbon has provided influential estimates of the SCC, using the three major models, DICE, PAGE and MERGE; see [Interagency Working Group on the Social Cost of Carbon \[2010\]](#), [Greenstone et al. \[2013\]](#), [Johnson and Hope \[2012\]](#) for details and the associated discussion.

and 90th percentiles (the latter two may be thought of as providing a “confidence bound” on the mean) along with those from the DICE-2016R2 using distributions suggested in Nordhaus [2018].

We provide a snapshot of the distribution of SCC over time in table 4 while table 5 details the effect of different sources of uncertainty and of adaptation on the SCC. From table 4, it is clear that the SCC increases steeply over time, possibly due to the non-linearity of climate-change-related damages. Furthermore, as is evident from the 10th and 90th percentiles, the spread of the social costs of carbon is large and increasing over time, where (in 2150) the 90th percentile is more than 20 times the 10th (from about eight times in 2020). This may be related to the fact that many parameter combinations, e.g. states with larger values of the damage exponent and climate sensitivity, will lead to large damages, leading to larger social costs, in the future. This is also the finding of e.g. Ackerman et al. [2010], which finds a large reduction in consumption resulting from large climate-change-related damages (for cases where climate sensitivity parameter is larger than 3 and the damage exponent is larger than 4).

| | | 2020 | 2030 | 2040 | 2050 | 2080 | 2100 | 2150 |
|--------------------|-----------------------------|-------------------------|------|------|------|------|------|------|
| AD-DICE2016 | | Full uncertainty | | | | | | |
| | 10 th percentile | 15 | 20 | 26 | 33 | 56 | 78 | 128 |
| | Mean | 62 | 87 | 116 | 151 | 310 | 474 | 1174 |
| | 90 th percentile | 136 | 189 | 259 | 336 | 677 | 1047 | 2842 |
| | | No uncertainty | | | | | | |
| | | 52 | 72 | 95 | 122 | 235 | 340 | 722 |
| DICE 2016R2 | No uncertainty | 35 | 49 | 67 | 88 | 176 | 263 | 403 |
| | Mean (MC) | 37 | 51 | 68 | 90 | 202 | 342 | 636 |

Table 4: Social costs of carbon in 2005 US\$

Our estimate of the SCC without uncertainty for the base period (2020) is comparable to those of many other modeling studies³³: it is higher than the \$42 implied in Nordhaus [2017], the \$37 in Nordhaus [2018]³⁴ the \$16 in Gillingham et al. [2018] and the \$20 in Traeger [2014] while lower than the \$125 in Crost and Traeger [2013] (figure 4). Similarly, our estimate for the longer term future, at year 2100, corresponds well with the figure in Nordhaus [2018], is similar to the range of values in Crost and Traeger [2013] (450-550, in figure 4), while higher than the simpler calibrated model in Traeger [2014], at about \$200.

Examining table 4, it is clear that the effect of uncertainty on the SCC is sizeable, where the mean value in 2020 is almost 20% higher than that without uncertainty. The difference between the two estimates (with and without uncertainty) slowly increases over time, and in 2150 the mean SCC with uncertainty is larger than that without uncertainty by 60%. The difference in mean SCC with and without uncertainty in our model is more than twice that in Nordhaus [2018], which is based on DICE-2016R2.³⁵

³³ It is worth noting that difference base years and discount rates make a strict comparison of the SCC across studies tricky. We therefore limit ourselves to broad comparisons only, and provide computations of the SCC for DICE-2016.

³⁴ This figure differs from that in AD-DICE2016 due to the climate sensitivity parameter being higher in AD-DICE2016.

³⁵ Note that in Nordhaus [2018] the SCC is calculated based on an optimal mitigation path in the absence of uncertainty, whereas our SCC estimated are based on an mitigation path optimised over uncertainty.

In summary, our estimates of the mean SCC, while comparable to those in the current literature, suggest that uncertainty does exert a significant effect on the SCC, and that an optimal carbon tax is likely increasing significantly over time (which is also supported by the increasing variability of the SCC over time). Further, our model suggests that, while following a similar path to that in DICE-2016R2 (starting moderate and increasing rapidly), increases in the SCC are much stronger.

As to the effect of different sources of uncertainty, table 5 suggests that damage-related uncertainty exerts a sizeable effect on the SCC, followed by climate uncertainty. These findings reinforce those from previous sections, where mitigation response was seen to be significantly sensitive to these two sources of uncertainty (see section 4.3). Other sources of uncertainty, together or in isolation, do not exert as much of an effect (and are therefore not presented in the tables). Finally, our estimates for the SCC when adaptation is excluded is almost twice as large as when it is included. This suggests that adaptation is an exceedingly important channel for welfare improvements, and is again a reflection of our findings previously, that mitigation was higher without adaptation than with. This also reflects the need to incorporate adaptation into estimates of the SCC.

| | 2020 | 2030 | 2040 | 2050 | 2080 | 2100 | 2150 |
|---------------------------------------|------------|------------|------------|------------|-------------|-------------|-------------|
| Full Uncertainty | | | | | | | |
| 10th percentile | 15 | 20 | 26 | 33 | 56 | 78 | 128 |
| Mean | <i>62</i> | <i>87</i> | <i>116</i> | <i>151</i> | <i>310</i> | <i>474</i> | <i>1174</i> |
| 90th percentile | 136 | 189 | 259 | 336 | 677 | 1047 | 2842 |
| Full Uncertainty–no adaptation | | | | | | | |
| 10th percentile | 20 | 27 | 36 | 46 | 85 | 114 | 199 |
| Mean | <i>171</i> | <i>249</i> | <i>346</i> | <i>467</i> | <i>1004</i> | <i>1502</i> | <i>3309</i> |
| 90th percentile | 396 | 579 | 801 | 1042 | 2288 | 3448 | 8314 |
| Damage uncertainty only | | | | | | | |
| 10th percentile | 23 | 30 | 39 | 49 | 91 | 133 | 284 |
| Mean | <i>61</i> | <i>86</i> | <i>116</i> | <i>150</i> | <i>296</i> | <i>429</i> | <i>872</i> |
| 90th percentile | 116 | 166 | 225 | 296 | 590 | 854 | 1698 |
| Climate uncertainty only | | | | | | | |
| 10th percentile | 19 | 26 | 34 | 43 | 81 | 116 | 232 |
| Mean | <i>51</i> | <i>71</i> | <i>94</i> | <i>121</i> | <i>236</i> | <i>345</i> | <i>721</i> |
| 90th percentile | 89 | 124 | 165 | 214 | 422 | 622 | 1308 |

Table 5: Social costs of carbon (2005 US\$): Sources of Uncertainty

5 Conclusions

The main aim of this paper has been to overcome some of the challenges related to using IAMs for climate policy analysis. To this end we develop a new IAM, AD-DICE2016 (based on DICE-2016R2), with which we demonstrate that IAMs can not only be used to determining robust optimal policy paths under extreme

(fat-tailed) uncertainty regarding model parameters but can also be simultaneously enriched to include an hitherto neglected yet important aspect, adaptation, and without any restrictive model simplifications.

The literature concerning optimal climate policy with uncertainty has thus far been largely addressed using either the Monte Carlo (MC) approach or the recursive Stochastic Dynamic Programming (SDP)-based approach. In the MC approach, using a large number of draws from parameter distribution, the optimal paths resulting from each draw is averaged to yield the average optimal path. This approach has been shown (here and in previous studies) to lead to incoherent policy conclusions. The SDP approach, which provides a robust and well-developed approach to analysing policies under uncertainty, has contributed significantly to understanding many dimensions of climate policy making under uncertainty. This approach, however, often necessitates model dimension reduction, in view of inherent computational challenges.

In this paper, we illustrate a method to account for parameter uncertainty using a state-contingent approach, wherein parameters take outcomes in a large yet finite number of states-of-the-world, and the decision maker optimizes over four control variables taking into account different outcomes in each state of the world. In contrast to the SDP approach, we are able to increase the state space dimension (to seven, from four in previous SDP versions of DICE) and the number of control variables (from two to four). Our non-recursive yet coherent approach to uncertainty is suited to contexts where parameters are uncertain (as in many IAMs), and, differs from MC approaches in that it considers every state-of-the-world when determining optimal policy. Unlike the SDP-based approach, however, it is not suited for contexts where the state variables evolve over time. Consequently, our approach to understanding the effect of parameter uncertainty on climate policy in IAMs complements SDP-based approaches to analysing the effects of structural uncertainty.

The AD-DICE2016 model includes several other important developments. Firstly, our model incorporates uncertainty over a large number of parameters (eight), each representing core aspects of the climate-economy causal chain, in contrast to common practice in the literature that examines uncertainty over very small number of parameters. Secondly, we incorporate more appropriate probability distributions based upon the most-recent knowledge, including “fat-tailed” distributions where recommended, allowing for extreme damage possibilities (of up to 99.8% of GDP). Thirdly, we choose to use a bounded burr utility function, following recommendations in both the theoretical and climate policy literature. Finally, to enable climate policy setting that reflects real-world contexts, we include, in addition to mitigation, two types of adaptation policies, reactive and proactive adaptation. This aspect is rather important in view of the differential time pattern of their costs and benefits, meaning that adaptation and mitigation policies interact in non-trivial ways.

Our model yields three broad insights regarding climate policy: The first relates to the importance of different sources of uncertainty: damage and climate uncertainty affect mitigation policies the most, whereas adaptation policies are greatly influenced by uncertainty over damages. Our findings broadly suggest that climate policy under uncertainty turns more aggressive i.e. mitigation is shifted forward in time and adaptation investments and spending are increased. Secondly, we find that adaptation and

mitigation policies under uncertainty exhibit strong inter-linkages, accentuating the need to consider both policies simultaneously. Third, we find that the SCC rises steeply over time, consistent with the findings of many recent deterministic studies. We also find that adaptation plays a significant role in the level of SCC. Moreover, with uncertainty, we find that the SCC increases faster over time, more so than in most studies in the extant literature. Such findings in the current IAM models have thus far been limited to cases where “catastrophes” (e.g. Ackerman et al. [2010], Dietz [2011]) were allowed for or controversial choices of discount rates or other key parameters were made.

More broadly, our study is among the many recent ones e.g. Traeger [2014], Cai et al. [2012], Crost and Traeger [2013], Ikefuji et al. [2020] which attempt to ameliorate many of the well known drawbacks of IAMs, in particular dealing with uncertainty and climate policy, each in a different way. The overarching message of these disparate modeling efforts, we believe, is this: that while the current crop of IAMs are unsatisfactory, suffering from an under-estimation of many risks (as pointed out in e.g. Stern [2013], Pindyck [2013]) and from ad-hoc functional forms (as forcefully argued in Pindyck [2013]), there are many alternative ways of improving these models. In particular, these modeling efforts illustrate a scope for optimism, that IAMs can be improved as needed for policy analysis and that the drawbacks of certain approaches to dealing with key aspects (such as the Monte Carlo approach to uncertainty, highlighted in Pindyck [2017], Crost and Traeger [2013]) of climate-economy frameworks need not mean that these frameworks are entirely unsuited for policy evaluation. In any case, alternative frameworks that attempt to quantify uncertain aspects of models, such as expert elicitation, suffer from their own drawbacks (Oppenheimer et al. [2016]). Consequently, the use of IAMs to “organise our necessarily incomplete perceptions” and “to make rough quantitative judgements about the consequences of economic policy..” (in the words of Solow [1985]) may be both necessary and, as recent modeling advances (including ours here) illustrates, perfectly feasible.

References

- F. Ackerman and E.A. Stanton. Climate Risks and Carbon Prices: Revising the Social Cost of Carbon. *Economics: The Open-Access, Open-Assessment E-Journal*, 6, 2012. <http://dx.doi.org/10.5018/economics-ejournal.ja.2012-10>.
- F. Ackerman, E.A. Stanton, and R. Bueno. Fat tails, exponents, extreme uncertainty: Simulating catastrophe in dice. *Ecological Economics*, 69(8):1657–1665, 2010.
- Barry Anderson, Emanuele Borgonovo, Marzio Galeotti, and Roberto Roson. Uncertainty in climate change modeling: Can global sensitivity analysis be of help? *Risk Analysis*, 34(2):271–293, 2014. <https://doi.org/10.1111/risa.12117>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/risa.12117>.

- D. Anthoff and R.S.J. Tol. The uncertainty about the social cost of carbon: A decomposition analysis using fund. *Climate Change*, 117:515–530, 2013. 10.1007/s10584-013-0706-7.
- D. Anthoff and R.S.J. Tol. Climate policy under fat-tailed risk: an application of FUND. *Ann Oper Res*, 220:223–237, 2014. 10.1007/s10479-013-1343-2.
- David Anthoff, Richard SJ Tol, and Gary W Yohe. Risk aversion, time preference, and the social cost of carbon. *Environmental Research Letters*, 4(2):024002, 2009.
- O. Bahn, K.C. de Bruin, and C. Fertel. Will Adaptation Delay the Transition to Clean Energy Systems? An Analysis with AD-MERGE. *The Energy Journal*, 0(Number 4), 2019. 10.5547/01956574.40.4.obah. URL <https://ideas.repec.org/a/aen/journal/ej40-4-debruin.html>.
- F. Bosello, C. Carraro, and E. De Cian. Climate policy and the optimal balance between mitigation, adaptation and unavoided damage. *Climate Change Economics*, 1(2), 2010.
- Valentina Bosetti, Carlo Carraro, Marzio Galeotti, Emanuele Massetti, and Massimo Tavoni. Witch a world induced technical change hybrid model. *The Energy Journal*, 27:13–37, 2006.
- M. Burke, Hsiang S.M., and Miguel E. Global non-linear effect of temperature on economic production. *Nature*, 527:235–239, 2015. <https://doi.org/10.1038/nature15725>.
- Y. Cai, K. Judd, and T. Lontzek. Dsice: A dynamic stochastic integrated model of climate and economy. RDCEP Working Paper 12-02, Center for Robust Decisionmaking on Climate & Energy Policy (RDCEP), 2012.
- Yongyang Cai, Kenneth L Judd, and Thomas S Lontzek. The social cost of carbon with economic and climate risks. *arXiv preprint arXiv:1504.06909*, 2015.
- Benjamin Crost and Christian P Traeger. Optimal climate policy: uncertainty versus monte carlo. *Economics Letters*, 120(3):552–558, 2013.
- K. C. de Bruin, R. B. Dellink, and S. Agrawala. *Economic Aspects of Adaptation to Climate Change: Integrated Assessment Modelling of Adaptation Costs and Benefits*, volume 6 of *OECD Working Papers*. OECD Publishing Paris, 2009b.
- K.C. de Bruin. Distinguishing between proactive (stock) and reactive (flow) adaptation. *CERE Working paper series*, 3(1), 2011.
- K.C. de Bruin and R.B. Dellink. How harmful are restrictions on adapting to climate change? *Global Environmental Change*, 21(1):34–45, 2011.
- K.C. de Bruin, R.B. Dellink, and R.S.J. Tol. Ad-dice: an implementation of adaptation in the dice model. *Climatic Change*, 95(1):63–81, 2009a.

- Simon Dietz. High impact, low probability? an empirical analysis of risk in the economics of climate change. *Climatic Change*, 108(3):519–541, 2011.
- Simon Dietz and Geir B Asheim. Climate policy under sustainable discounted utilitarianism. *Journal of environmental economics and management*, 63(3):321–335, 2012.
- Simon Dietz, Christian Gollier, and Louise Kessler. The climate beta. *Journal of Environmental Economics and Management*, 2017.
- R.S.J. Fankhauser, S. Tol. On climate change and economic growth. *Resource and Energy Economics*, 27(1):1–17, 2005.
- T. Felgenhauer and K.C. de Bruin. The optimal paths of climate change mitigation and adaptation under certainty and uncertainty. *International Journal of Global Warming*, 1(1):66–88, 2009.
- Reyer Gerlagh and Matti Liski. Carbon prices for the next hundred years. *The Economic Journal*, 2016.
- K. Gillingham, W. Nordhaus, D. Anthoff, G. Blanford, V. Bosetti, P. Christensen, H. McJeon, and J. Reilly. Modeling uncertainty in integrated assessment of climate change: A multimodel comparison. *Journal of the Association of Environmental and Resource Economists*, 5(4), 2018.
- N. Glanemann, S.N. Willner, and A. Levermann. Paris climate agreement passes the cost-benefit test. *Nature Communication*, 11(1), 2020.
- Michael Greenstone, Elizabeth Kopits, and Ann Wolverton. Developing a social cost of carbon for us regulatory analysis: A methodology and interpretation. *Review of Environmental Economics and Policy*, 7(1):23–46, 2013.
- Michael Hanemann. What is the economic cost of climate change? In Stephen Schneider, Armin Rosencranz, Michael Mastrandrea, and Kristin Kuntz-Duriseti, editors, *Climate Change Science and Policy*, pages 185–193. Island Press, 2009.
- G. Heal. The Economics of the Climate. *Journal of Economic Literature*, 55(3):1046–1063, 2017. <https://doi.org/10.1257/jel.20151335>.
- A.F. Hof, M.G.J. den Elzen, and D.P. van Vuuren. Analysing the costs and benefits of climate policy: Value judgements and scientific uncertainties. *Global Environmental Change*, 18:412–424, 2008.
- Chris Hope. The marginal impact of co2 from page2002: An integrated assessment model incorporating the ipcc’s five reasons for concern. *Integrated assessment*, 6(1), 2006.
- M. Ikefuji, R. Laeven, J.R. Magnus, and C. Muris. Weitzman meets nordhaus: Expected utility and catastrophic risk in a stochastic economy-climate model. Technical Report Discussion Paper 825, Osaka University, ISER, 2011.

- M. Ikefuji, R.J.A. Laeven, J.R. Magnus, and C. Muris. Pareto utility. *Theory and Decision*, pages 1–15, 2013.
- M. Ikefuji, R.J.A. Laeven, J.R. Magnus, and C. Muris. Expected utility and catastrophic risk in a stochastic economy–climate model. *Journal of Econometrics*, 214(1), 2020.
- A. Ingham, J. Ma, and A. Ulph. Climate change, mitigation and adaptation with uncertainty and learning. *Energy Policy*, 35(11):5354–5369, 2007.
- Interagency Working Group on the Social Cost of Carbon. Social cost of carbon for regulatory impact analysis under executive order 12866. Technical report, White House, Genf, Schweiz, 2010. URL <http://www.whitehouse.gov/sites/default/files/omb/inforeg/for-agencies/Social-Cost-of-Carbonfor-RIA.pdf>.
- Matthew O Jackson and Leeat Yariv. Collective dynamic choice: the necessity of time inconsistency. *American Economic Journal: Microeconomics*, 7(4):150–178, 2015.
- Laurie T Johnson and Chris Hope. The social cost of carbon in us regulatory impact analyses: an introduction and critique. *Journal of Environmental Studies and Sciences*, 2(3):205–221, 2012.
- N. L. Johnson, S. Kotz, and N. Balakrishnan. *Continuous univariate distributions*. Wiley & Sons, 1995.
- Larry Karp and Jiangfeng Zhang. Regulation with anticipated learning about environmental damages. *Journal of Environmental Economics and Management*, 51(3):259–279, 2006.
- S Kotz and JR Van Dorp. *Beyond Beta, Other Continuous Families of Distributions with Bounded Support and Applications*. World Scientific Press, Republic of Singapore, 2004.
- Derek Lemoine. The climate risk premium: how uncertainty affects the social cost of carbon. *Journal of the Association of Environmental and Resource Economists*, 8(1):27–57, 2021.
- Alan Manne, Robert Mendelsohn, and Richard Richels. Merge: A model for evaluating regional and global effects of ghg reduction policies. *Energy Policy*, 23(1):17–34, 1995.
- G. A. Meehl and T. F. et al Stocker. Global climate projections. climate change 2007: The physical science basis. contribution of working group i to the fourth assessment report of the intergovernmental panel on climate change. *Review of Environmental Economics and Policy*, 2007.
- Antony Millner. On welfare frameworks and catastrophic climate risks. *Journal of Environmental Economics and Management*, 65(2):310–325, 2013.
- M.G. Morgan and H. Dowlatabadi. Learning from integrated assessment of climate change. *Climatic Change*, 34(3):337–368, 1996.

- W.D. Nordhaus. *A question of balance: Weighing the options on global warming policies*. Yale Univ Pr, 2008.
- William Nordhaus. Projections and uncertainties about climate change in an era of minimal climate policies. *American Economic Journal: Economic Policy*, 10(3):333–60, August 2018. 10.1257/pol.20170046. URL <https://www.aeaweb.org/articles?id=10.1257/pol.20170046>.
- William D Nordhaus. The economics of tail events with an application to climate change. *Review of Environmental Economics and Policy*, 5(2):240–257, 2011.
- William D Nordhaus. Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, page 201609244, 2017.
- Roman Olson, Ryan Sriver, Marlos Goes, Nathan M. Urban, H. Damon Matthews, Murali Haran, and Klaus Keller. A climate sensitivity estimate using bayesian fusion of instrumental observations and an earth system model. *Journal of Geophysical Research: Atmospheres*, 117(D4), 2012.
- Michael Oppenheimer, Christopher M Little, and Roger M Cooke. Expert judgement and uncertainty quantification for climate change. *Nature climate change*, 6(5):445–451, 2016.
- S. Peterson. Uncertainty and economic analysis of climate change: A survey of approaches and findings. *Environmental Modeling and Assessment*, 11(1):1–17, 2006.
- Robert S Pindyck. Climate change policy: What do the models tell us? *Journal of Economic Literature*, 51(3):860–872, 2013.
- Robert S Pindyck. The use and misuse of models for climate policy. *Review of Environmental Economics and Policy*, 11(1):100–114, 2017.
- F. Piontek, M. Kalkuhl, E. Kriegler, A. Schultes, M. Leimbach, O. Edenhofer, and N. Bauer. Economic Growth Effects of Alternative Climate Change Impact Channels in Economic Modeling. *Environmental and Resource Economics*, 73:1357–1385, 2019. <https://doi.org/10.1007/s10640-018-00306-7>.
- W.A. Pizer. The optimal choice of climate change policy in the presence of uncertainty. *Resource and Energy Economics*, 21(3-4):255–287, 1999.
- Richard G Richels, Alan S Manne, and Tom ML Wigley. Moving beyond concentrations: the challenge of limiting temperature change. Working Paper 04-11, AEI-Brookings Joint Cente, 2004.
- T. Roughgarden and S.H. Schneider. Climate change policy: quantifying uncertainties for damages and optimal carbon taxes. *Energy Policy*, 27(7):415–429, 1999. [https://doi.org/10.1016/S0301-4215\(99\)00030-0](https://doi.org/10.1016/S0301-4215(99)00030-0).

- Barry Smit, O. V. Pilifosova, I. Burton, B. Challenger, S. Huq, Richard J. T. Klein, and Gary W. Yohe. Adaptation to climate change in the context of sustainable development and equity. In James J. McCarthy, Osvaldo F. Canziani, Neil A. Leary, David J. Dokken, and Katherina S. White, editors, *Climate Change 2001: Impacts, Adaptation, and Vulnerability: Contribution of Working Group II to the Third Assessment Report of the Intergovernmental Panel on Climate Change*, pages 877–911. Press Syndicate of the University of Cambridge, Cambridge UK, 2001.
- Robert M Solow. Economic history and economics. *The American Economic Review*, 75(2):328–331, 1985.
- Nicholas Stern. The structure of economic modeling of the potential impacts of climate change: grafting gross underestimation of risk onto already narrow science models. *Journal of Economic Literature*, 51(3):838–859, 2013.
- R.S.J. Tol. Why worry about climate change? a research agenda. *Environmental Values*, 17(4):437–470, 2008.
- R.S.J. Tol. The Economic Impacts of Climate Change. *Review of Environmental Economics and Policy*, 12(1):4–25, 2018. 10.1093/reep/rex027.
- Christian Traeger. A 4-stated dice: Quantitatively addressing uncertainty effects in climate change. *Environmental and Resource Economics*, 59:1–37, 2014.
- United Nations, Department of Economic and Social Affairs, Population Division. World population prospects: the 2017 revision, key findings and advance tables. Technical Report Working Paper ESA/P/WP/248, Department of Economics and Social Affairs, Population Division, 2017.
- M.B.A. van Asselt and J. Rotmans. Uncertainty in integrated assessment modelling. *Climatic Change*, 54(1):75–105, 2002.
- M.L. Weitzman. On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics*, 91(1):1–19, 2009.
- M.L. Weitzman. Fat-tailed uncertainty in the economics of catastrophic climate change. *Review of Environmental Economics and Policy*, 5(2):275–292, 2011.
- M.L. Weitzman. Ghg targets as insurance against catastrophic climate damages. *Journal of Public Economic Theory*, 14(2):221–244, 2012.
- John Weyant. Some contributions of integrated assessment models of global climate change. *Review of Environmental Economics and Policy*, 11(1):115–137, 2017.
- I.S. Wing and E. Lanzi. Integrated assessment of climate change impacts: Conceptual frameworks, modelling approaches and research needs. Working Paper 66, OECD, 2014.

G.W. Yohe, R.D. Lasco, Q.K. Ahmad, N.W. Arnell, S.J. Cohen, C. Hope, A.C. Janetos, and R.T. Perez. Perspectives on climate change and sustainability. In M.L. Parry, O.F. Canziani, J.P. Palutikof, P.J. van der Linden, and C.E. Hanson., editors, *Climate Change 2007: Impacts, Adaptation and Vulnerability: Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, pages 811–841. Cambridge University Press, 2007.

Kirsten Zickfeld, M. Granger Morgan, David J. Frame, and Keith David W. Expert judgements about transient climate response to alternative future trajectories of radiative forcing. *Proceedings of the National Academy of Sciences*, 106:16129–16135, 2010.

Stéphane Zuber. The aggregation of preferences: can we ignore the past? *Theory and Decision*, 70(3): 367–384, 2011.

Appendix A Static Welfare Framework

We introduce some notions of welfare, and its associated notation, that will help fix the precise idea of preference aggregation necessitated when preference parameters vary. Denote by X the set of all alternative states (called “social states” in the social choice literature), by $N := \{1, 2, \dots, n\}$ the number of “individuals” in an economy, with each individual’s preferences represented by a utility function U_i (clarified further below), defined on X . Uncertainty can be introduced by identifying X to be set of all probability distributions L defined on a finite set of basic outcomes $B := \{b_1, \dots, b_J\}$ i.e. with each lottery x is associated a vector $\pi = (\pi_1, \dots, \pi_J)$ of probabilities of obtaining outcomes $b = (b_1, \dots, b_J)$; U_i is interpreted as individual i ’s VonNeuman-Morgenstern (vNM) Utility function.

Let $W : X \rightarrow \mathbb{R}$ be a social welfare function (SWF) and let welfare be represented by an expected social welfare function i.e. social preferences on X (alternatively, on lotteries L) can be represented by the expected value of a “vNM-Bergson”-type SWF, W . Then, Harsanyi’s Social Aggregation Theorem has the implication that the social welfare functional of society is ordinally equivalent to (with $\{\omega_i\}$ representing welfare weights)


$$W(x) = \sum_{i=1}^N \omega_i U_i(x) = \sum_{i=1}^N \omega_i \mathbb{E}[u_i(b_j)], \quad \omega_i > 0, \quad (\text{A.1})$$

which is of the weighted utilitarian form commonly used. Essentially, Harsanyi’s theorem provides a way to aggregate preferences of individuals each of whom is an expected utility maximizer. In keeping with the expected utility approach, society too is an expected utility maximizer and this results in the weighted utilitarian manner of evaluating social prospects in (A.1).


Note that the typical preference aggregation is *over individuals* in a society who have different preferences, **not** over the same individuals with “uncertain” preferences. The interpretation of eq. (17) in [Pizer \[1999\]](#) when aggregating over “preferences states of nature” is unclear, and there is no discussion of this issue in that study. However, we argue that it is possible to reinterpret the setting as follows. Let there be M “types” of individuals in every given population³⁶; these individuals differ from one another in the parameters of their utility function and their evaluation of inter-temporal trade-offs (i.e. pure rate of time preference). In this case, aggregation over these M “individual types” have a straightforward interpretation as a *social aggregate* over the different types of individuals, and one is back in the Harsanyi-like setting above.

³⁶ The individual types do not correspond to the population size, and are to be thought of as independent of it and invariant over time, unlike population. Indeed, so far, our discussion has completely ignored the issue of population (implicitly assuming a fixed population) while in our model, population is one of the uncertain parameters. We do not introduce population-ethics based social choice since it is not clear that the added complication in our case can yield additional insights. See [Dietz and Asheim \[2012\]](#) (and references therein) for a brief discussion of the issue of social welfare for varying populations in the context of climate change.

Appendix B Additional Figures



(a) Effect of varying κ on utility function.



(b) Effect of varying λ on RRA.

Figure B.1: Effects of κ and λ on Burr utility function and its coefficient of RRA.




Figure B.2: Comparing CRRA, Burr and Exponential utility functions (at low consumption values).

Appendix C Equations of AD-DICE-2016

Here we present the primary model equations for completeness. We note that certain equations used to calculate parameter values over time have been omitted unless they directly concern the uncertain parameters. These parameters are identical to those in DICE-2016R2, and readers are directed to Nordhaus [2017] and Nordhaus [2018] for a fuller description of these components of the model. Time is indexed by t , the planning period by $T = 1, 2, \dots, T$, and the (uncertain) states of the world are indexed by s , $S = 1, 2, \dots, S$. In what follows, parameters are written in lower case letters and variables in upper case (except for Greek letters, where we do not make this distinction).

The model maximises utility subject to 7 state variable (K , SAD , M_{AT} , M_{UP} , M_{LO} , T_{AT} , T_{LO}) and 4 control variables (MU , I , IAD , FAD). The following 10 parameters are drawn randomly from probability distributions: κ_s , pop_s , $ga_{1,s}$, $gsig_{1,s}$, $fosl_{1,s}$, $\phi_{12,s}$, $\alpha_{3,s}$, $t2xco2_s$.

C.1 Utility

Utility is given by U :

$$U = \frac{1}{N} \sum_{s=1}^N \left(\sum_{t=1}^T R^t L_{t,s} P U_{t,s} \right) \quad (C.1)$$

$P U_{t,s}$ represents the utility per period and state and is computed using the burr utility function,

$$P U_{t,s} = 1 - \left(\frac{\lambda_s}{\frac{C_{t,s}}{L_{t,s}} + \lambda_s} \right)^{\kappa_s}, \quad (C.2)$$

with λ_s and κ_s the Burr utility-related parameters, $C_{t,s}$ total consumption and $L_{t,s}$ population. β^t is the discount factor computed as $R^t = (1 + \rho_t)^{-t}$

C.2 Output

Output before damages and abatement is determined by a Cobb-Douglas function of capital (K) and labour (L). Where $al_{t,s}$ represents the total factor productivity per region per time period:

$$Y_{GROSS_{t,s}} = al_{t,s} K_t^\beta L_{t,s}^{1-\beta} \quad (C.3)$$

$$L_{t+1,s} = L_{t,s} (pop_s / L_{t,s})^{popad} \quad (C.4)$$

Total factor productivity is determined over time as a function of $ga_{1,s}$:

$$al(t, s) = al(1, s) / ((1 - ga(t, s))^t) \quad (C.5)$$

$$ga_{t,s} = ga_{1,s}^{-delta_{ga} * t} \quad (C.6)$$

Net output is given as the gross output minus net climate change damages (residual damages and adaptation) and mitigation costs:

$$Y_{NET,t,s} = Y_{GROSS,t,s} (1 - RD_{t,s} - PC_t - MC_t). \quad (C.7)$$

The consumption function is given by:

$$C_{t,s} = Y_{NET,t,s} - I_t. \quad (C.8)$$

Capital accumulation is defined in the conventional manner:

$$K_{t+1} = (1 - \delta_k) K_t + I_t, \quad (C.9)$$

Where δ_k is the depreciation rate and I_t the investments in capital.

C.3 Emissions and mitigation

In the DICE-2016 and AD-DICE2016 models, mitigation—modelled as a fraction of total emissions—represents a control or policy variable which reduces GHG emissions per unit of production. The relationship between economic production and GHG emissions is described by the parameter σ_t , assumed to exogenously decrease over time (reflecting the decarbonization of production via technological advances). Industrial emissions (E_{IND}) are determined by the level of production, the value of σ_t and by (the control variable of) mitigation (MU_t) as

$$E_{IND,t,s} = \sigma_{t,s} Y_{GROSS,t,s} (1 - MU_t), \quad (C.10)$$

where $\sigma, \mu \in [0, 1]$.

In the DICE-2016 and AD-DICE2016 models, σ_t is set at 0.13 in 2005 and declines at a rate of 0.3% per decade thereafter;

$$\sigma_{t+1,s} = \sigma_{t,s}^{gsig} \quad (C.11)$$

$$gsig_{t+1,s} = gsig_{t,s}^{1+dsig} \quad (C.12)$$

Naturally mitigation is carried out at a cost, determined as

$$MC_t = \theta_{1,t} \mu_t^{\theta_{2,t}}, \quad (C.13)$$

where $\theta_1, \theta_2 > 0$. In the DICE-2016 and AD-DICE2016 models θ_1 depends on the backstop technology price and the level of carbon per unit of output, and is set at 0.2 initially (and decreases over time) while θ_2 is set at 2.8. The mitigation levels and costs represent an aggregate of possible mitigation options and are updated in each version of DICE. Carbon sinks are included in the model after 2150 and are limited to 20% of greenhouse gas emissions, hence in periods after 2150 emissions can be negative when full mitigation is combined with carbon sinks.

Total emissions are given as the total of industrial emissions and emissions from land use

$$E_{t,s} = E_{IND,t,s} + E_{LU,t,s} \quad (C.14)$$

C.4 Damages and adaptation

Gross damages from climate change are given as a fraction of gross output as follows:

$$GD_{t,s} = \alpha_1 T_{AT,t,s} + \alpha_2 T_{AT,t,s}^{\alpha_3}, \quad (C.15)$$

where $\alpha_2 > 0$ and $\alpha_3 > 0$. These are the damages that occur if no adaptation takes place, and are thus higher than the net damages. These gross damages can be reduced through the use of adaptation to residual damages. Residual damages depend on both the gross damages (GD_t) and the total level of adaptation (PT_t), both stock and flow as follows:

$$RD_{t,s} = \frac{GD_{t,s}}{1 + PT_{t,s}}. \quad (C.16)$$

Both forms of adaptation are imperfect substitutes for each other, and are aggregated using a Constant Elasticity of Substitution (CES) function. Together with the residual damages function, this describes how adaptation expenditures reduce the damages caused by climate change:

$$PT_t = \gamma_1 \cdot (\beta \cdot SAD_t^\rho + (1 - \beta) \cdot FAD_t^\rho)^{\frac{\gamma_2}{\rho}}, \quad (C.17)$$

where SAD_t is the total amount of adaptation capital stock at time t and FAD_t is the amount spent on reactive adaptation in period t . Furthermore, $\rho = (\sigma - 1)/\sigma$ is the elasticity of substitution between stock and flow adaptation, with $\beta \in [0, 1]$, $\rho \in (-\infty, 1]$ and $\gamma_1, \gamma_2 > 0$. Adaptation capital stock is accumulated in the same manner as conventional capital stock:

$$SAD_{t+1} = (1 - \delta_k)SAD_t + IAD_t, \quad (C.18)$$

where δ_k is the depreciation rate and IAD_t are the investments in stock adaptation. The total adaptation costs in each period are thus

$$PC_t = FAD_t + IAD_t. \quad (C.19)$$

C.5 Climate cycle

The CO_2 stock in the atmosphere in each period is given as:

$$M_{AT,t+1,s} = \phi_{11,s}M_{AT,t,s} + \phi_{21,s}M_{UP,t,s} + E_{t,s}, \quad (C.20)$$

The CO_2 stock in the upper oceans is given as:

$$M_{UP,t+1,s} = \phi_{12,s}M_{AT,t,s} + \phi_{22,s}M_{UP,t,s} + \phi_{32,s}M_{LO,t,s}, \quad (C.21)$$

The CO_2 stock in the lower oceans is given as:

$$M_{LO,t+1,s} = \phi_{23,s}M_{UP,t,s} + \phi_{33,s}M_{LO,t,s}, \quad (C.22)$$

Radiative forcing is given as:

$$F_{t,s} = \eta \{ \log[M_{AT,t,s}/M_{AT,1750,s}] / \log 2 \} + F_{EX,t}, \quad (C.23)$$

Global mean atmospheric surface temperature:

$$T_{AT,t+1,s} = T_{AT,t,s} + \xi_1 \{F_{t,s} - (fco22x/t2xco2_s)T_{AT,t,s} - \xi_3 [T_{AT,t,s} - T_{LO,t,s}]\}, \quad (C.24)$$

Temperature in the lower oceans is given as:

$$T_{LO,t+1,s} = T_{LO,t,s} + \xi_4 \{T_{AT,t,s} - T_{LO,t,s}\}, \quad (C.25)$$

C.6 The Optimization Problem

The AD-DICE2016 model optimises by maximising utility (as defined in equation C.1) using the control variables (MU , I , IAD , FAD) and subject to 7 state variable (K (see eq. C.9), SAD (see eq. C.18), M_{AT} (see eq. C.20), M_{UP} (see eq. C.21), M_{LO} (see eq. C.22), T_{AT} (see eq. C.24), T_{LO} (see eq. C.25)) :

$$\begin{aligned} & \max_{I_t, MU_t, IAD_t, FAD_t} U \\ & s.t. K_{t+1} = f(K_t, I_t) \\ & SAD_{t+1} = (1 - \delta_k)SAD_t + IAD_t \\ & M_{AT,t+1,s} = f(M_{AT,t,s}, M_{UP,t,s}, I_t, MU_t, \theta) \\ & M_{UP,t+1,s} = \phi_{12,s}M_{AT,t,s} + \phi_{22,s}M_{UP,t,s} + \phi_{32,s}M_{LO,t,s} \\ & M_{LO,t+1,s} = \phi_{23,s}M_{UP,t,s} + \phi_{33,s}M_{LO,t,s} \\ & T_{AT,t+1,s} = T_{AT,t,s} + \xi_1 \{F_{t,s} - (fco22x/t2xco2_s)T_{AT,t,s} - \xi_3 [T_{AT,t,s} - T_{LO,t,s}]\} \\ & T_{LO,t+1,s} = T_{LO,t,s} + \xi_4 \{T_{AT,t,s} - T_{LO,t,s}\} \end{aligned} \quad (C.26)$$