

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Uncertainty, Inequality, and Global Climate Policy

Permalink

<https://escholarship.org/uc/item/6kf949x0>

Author

Errickson, Frank Cooper

Publication Date

2020

Peer reviewed|Thesis/dissertation

Uncertainty, Inequality, and Global Climate Policy

By

Frank C. Errickson

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Energy and Resources

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor David Anthoff, Chair

Professor William D. Collins

Professor Klaus Keller

Summer 2020

Uncertainty, Inequality, and Global Climate Policy

Copyright 2020

by

Frank C. Errickson

Abstract

Uncertainty, Inequality, and Global Climate Policy

by

Frank C. Errickson

Doctor of Philosophy in Energy and Resources

University of California, Berkeley

Professor David Anthoff, Chair

Despite decades of research on the science and potential impacts of climate change, considerable disagreement remains over how to balance the upfront costs of reducing greenhouse gas emissions with the benefits of avoiding future negative impacts. This disagreement persists, in part, because deep uncertainties in the climate system limit our understanding of how temperature and sea level will evolve over the coming centuries. The global nature and long timescales of climate change provide an additional challenge, as climate policy analysis requires difficult ethical judgments on how to value the well-being of individuals at different income levels and across generations. In this dissertation, I address these issues by using integrated assessment models (IAMs) to analyze how distributional equity concerns and previously neglected climate uncertainties affect climate policy design. In the first part, I calculate the optimal carbon dioxide mitigation policy when society accounts for the trade-off between mitigation costs, climate damages, and the climate and localized health consequences of changes in air pollutant co-emissions. The presence of health “co-benefits” leads to increased mitigation levels that may be consistent with a 2 °C target, but the magnitude of this effect crucially depends on independent air quality policies and the value society assigns to improvements in human health. In the second part, I use a Bayesian calibration framework and an ensemble of IAMs to quantify how parametric climate uncertainties affect the social cost of methane (SC-CH₄). Despite accounting for the recent 25% upward revision to methane radiative forcing, the constrained models produce a mean SC-CH₄ estimate 23% lower than the value recently used by the U.S. federal government. I also provide the first probabilistic equity-weighted SC-CH₄ estimates and find that the U.S. value is more than an order of magnitude higher than sub-Saharan Africa’s. The third part of my dissertation focuses on low-probability, high-impact upper tail estimates of the social cost of carbon (SCC). In this analysis, I first show that the simple climate models from the IAMs used to calculate official U.S. SCC estimates do not pass relatively simple out-of-sample hindcast tests and exhibit impulse response behaviors that are inconsistent with current scientific understanding. I then couple the non-climate components of each IAM to a reduced complexity Earth system model and quantify how different treatments of parametric uncertainties affect 95th percentile SCC estimates. While failing to account for interactions between the uncertain parameters increases the 95th percentile SCC by more than 20%, following the official U.S. SCC climate uncertainty framework nearly doubles it.

For Hannah

TABLE OF CONTENTS

Table of Contents	ii
List of Figures	iv
List of Tables	vii
Acknowledgements	viii
1 Introduction	1
2 The Impact of Human Health Co-Benefits on Evaluations of Global Climate Policy	4
2.1 Introduction.....	4
2.2 Methods.....	5
2.3 Main Results	15
2.4 Independent Air Quality Control	19
2.5 Discounting.....	21
2.6 Monetizing and Valuing Health Benefits	22
2.7 Additional Sensitivities in RICE + AIR	23
2.8 Exploring Model Uncertainty by Linking AIR to the FUND Model	24
2.9 Discussion and Conclusion	25
3 Improving the Climate Model Calibration Reduces the Social Cost of Methane	28
3.1 Introduction.....	28
3.2 Methods.....	30
3.3 Observations Constrain SC-CH4 Estimates.....	39
3.4 Non-Linear Parameter Relationships.....	41
3.5 Model Structure and SC-CH4 Variation.....	43
3.6 Equity-Weighting.....	45
3.7 Discussion and Caveats.....	47
4 Enhanced Representation of Climate Model Uncertainties Reduces Upper-Tail Social Cost of Carbon Estimates	48
4.1 Introduction.....	48
4.2 Methods.....	50
4.3 Variation in Climate Hindcasts and Impulse Response Behaviors.....	60
4.4 Correlations Constrain Upper Tail SCC Estimates.....	64
4.5 Sensitivity Analyses and Discussion	66
4.6 Conclusion	67
5 Conclusion	69

A	Supplementary Information for Chapter 2	72
A.1	Supplementary Figures	72
A.2	Supplementary Tables.....	80
A.3	Supplementary Notes	83
B	Supplementary Information for Chapter 3	87
B.1	Supplementary Figures.....	87
B.2	Supplementary Tables	94
C	Supplementary Information for Chapter 4	97
C.1	Supplementary Figures.....	97
C.2	Supplementary Tables	108
	References	113

LIST OF FIGURES

Figure 2.1: Diagram of the AIR module	8
Figure 2.2: Optimal decarbonization and temperature	16
Figure 2.3: Costs and benefits of mitigation	17
Figure 2.4: Health benefits of carbon mitigation.....	18
Figure 2.5: Impact of independent air quality control.....	20
Figure 2.6: The influence of pure time preference	22
Figure 2.7: FUND + AIR results	25
Figure 3.1: Probabilistic climate hindcast skills vary with model choice and treatment of parametric uncertainties.....	40
Figure 3.2: SC-CH ₄ distribution estimates	41
Figure 3.3: Non-linear climate parameter relationships constrain probabilistic temperature projections and SC-CH ₄ estimates	42
Figure 3.4: Discounted climate damage impulse response behavior and uncertainties strongly differ between DICE and FUND	44
Figure 3.5: Variations in equity-weighted SC-CH ₄ estimates and their uncertainties for different consumption elasticity of marginal utility values.....	46
Figure 4.1: Annualized versions of DICE-Climate and PAGE-Climate produce similar output to the default versions.....	52
Figure 4.2: Shares of historic global anthropogenic sulfur dioxide emissions for the eight PAGE-Climate regions	53
Figure 4.3: Comparison of RCP carbon dioxide emissions with the emission scenarios used for official U.S. social cost of carbon estimates	54
Figure 4.4: Temperature pattern scaling in PAGE	58
Figure 4.5: Climate model structure and parameter correlations influence out-of- sample temperature hindcast skill.....	62
Figure 4.6: Temperature response dynamics and uncertainty levels resulting from a 100 GtC emission pulse vary with climate model parameter choice.....	63
Figure 4.7: Social cost of carbon uncertainty and 95th percentile estimates vary with treatment of parametric uncertainties	65

Figure 4.8: Climate model calibration and parameter correlations constrain high social cost of carbon estimates.....	66
Figure A.1: Model comparison: RICE2010, Reference case, and RICE+AIR	72
Figure A.2: Optimal carbon prices	73
Figure A.3: Optimal decarbonization rates with a 3% discount rate.....	74
Figure A.4: Emission intensities by region and pollutant over time	75
Figure A.5: Baseline (pre-mitigation) emissions by region	76
Figure A.6: Co-reduction relationships	77
Figure A.7: Radiative forcing.....	78
Figure A.8: Cost curves for air quality control	79
Figure B.1: Additional probabilistic climate hindcasts	87
Figure B.2: Additional SC-CH4 estimates under different model structures and treatments of parametric uncertainty	88
Figure B.3: Using the RCP2.6 scenario keeps expected temperature projections below 2 °C and reduces SC-CH4 estimates relative to RCP8.5	89
Figure B.4: Robustness of SC-CH4 distributions to wider prior assumptions, and posterior methane cycle parameter relationships.....	90
Figure B.5: Additional models showing non-linear climate parameter relationships constrain SC-CH4 estimates	91
Figure B.6: Neglecting relationships between posterior parameter estimates increases probabilistic temperature projection uncertainty.....	92
Figure B.7: Strong positive relationship between the SC-CH4 and ECS under the U.S. SC-CH4 estimation framework	93
Figure C.1: The DICE, FUND, and PAGE climate models exhibit limited temperature and sea level hindcast skill under their default parameterizations.....	97
Figure C.2: Posterior climate model parameter distributions	98
Figure C.3: Posterior parameter correlations emerge after calibrating the DICE-Climate, FUND-Climate, and PAGE-Climate models to past observations	99
Figure C.4: Climate model structure and parameter correlations influence out-of-sample sea level hindcast skill	100
Figure C.5: Forcing scenarios and different treatments of parametric uncertainty	

	influence a variety of model responses to a 100 GtC emissions pulse	101
Figure C.6:	Climate model parameter correlations reduce uncertainty in the temperature response to cumulative carbon dioxide emissions	102
Figure C.7:	Sea level response dynamics and uncertainty levels resulting from a 100 GtC emission pulse vary with climate model parameter choice	103
Figure C.8:	Social cost of carbon distributions under different discount rates	104
Figure C.9:	Temperature impulse response dynamics from a 100 GtC emission pulse under the RCP2.6 scenario	105
Figure C.10:	Climate model projections and social cost of carbon distributions differ markedly between the RCP2.6 and RCP8.5 scenarios	106
Figure C.11:	Equity-weighted social cost of carbon distributions for different geographic regions and treatments of parametric climate uncertainties	107

LIST OF TABLES

Table 2.1: Impact of independent air quality control.....	20
Table 2.2: Monetizing life-years/lives	23
Table 2.3: Additional sensitivity analyses	24
Table A.1: Effect of valuing life-years differently in wellbeing terms in the optimization	80
Table A.2: The modified objective function.....	81
Table A.3: Radiative forcing coefficients by region in 2050 and 2100.....	82
Table B.1: Past SC-CH4 estimates and their treatment of parametric climate uncertainty.....	94
Table B.2: Uncertain climate model and statistical process parameters with their prior distributions.....	95
Table B.3: Model and scenario specific SC-CH4 estimates	96
Table C.1: Uncertain parameter distributions for DICE-Climate’s default settings and the prior distributions used for the Bayesian calibration	108
Table C.2: Uncertain parameter distributions for FUND-Climate’s default settings and the prior distributions used for the Bayesian calibration	109
Table C.3: Uncertain parameter distributions for PAGE-Climate’s default settings and the prior distributions used for the Bayesian calibration	110
Table C.4: Uncertain parameters for SNEASY+BRICK and their prior distributions.....	111
Table C.5: Model and scenario specific social cost of carbon estimates.....	112

ACKNOWLEDGEMENTS

This dissertation is not an endeavor entirely of my own making. It involved a healthy dose of luck and, more importantly, the unconditional support of my friends, family, and advisors. The brevity of these acknowledgements does not signal a lack of gratitude, but rather my inability to express my full appreciation for the tremendous levels of help and support I have been given along the way. I truly could not have done it alone.

I thank my parents for inspiring a love of the natural world and giving me the courage to pursue my dreams, even when I could not see a way forward.

I thank my early teachers, not only for their dedication and patience when teaching me to read, but for also giving me that first spark of excitement that occurs when you discover something new.

I thank Annie for making love and laughter a daily part of my life and always seeing the best in me even when I could not.

I thank my advisors – David Anthoff, Klaus Keller, Bill Collins, and Marc Fleurbaey – for their never-ending support, seemingly infinite patience, and contagious level of excitement for the work we do. I accomplished more than I ever thought possible because of your help and guidance. I am forever grateful.

Lastly, I thank Hannah, who has been at my side every step of the way and showed me it is possible to fit more joy into a single lifetime than I ever thought possible.

CHAPTER 1

Introduction

Modeling based research over the past several decades has produced important insights into society's complex relationship with the global climate system (National Academies of Sciences, Engineering, and Medicine 2017). We now have a deeper understanding of the impacts that may arise from a changing climate as well as the potential costs and benefits of different mitigation scenarios (Moore and Diaz 2015; Hsiang et al. 2017; National Academies of Sciences, Engineering, and Medicine 2017; Burke, Davis, and Diffenbaugh 2018; Lamontagne et al. 2019). The implications of these and related works for climate risk management partially rest on the assumption that coupled socioeconomic and natural systems can be reasonably modeled a century or more into the future. Importantly, this includes an ability to quantify the key, decision-relevant sources of uncertainty. When researchers have a mechanistic understanding of the underlying processes being modeled, this assumption may be quite reasonable. However, the complex nature of the climate problem and the presence of deep uncertainties oftentimes challenges our ability to make probabilistic statements about the future (Lempert et al. 2006). In the climate system, examples of such deep uncertainties include the risk of accelerated sea level from a collapse of the West Antarctic ice sheet, the temperature threshold at which a rapid dieback of the Amazon rainforest occurs, and imprecise knowledge about how much warming society will actually experience from a doubling of atmospheric carbon dioxide (CO₂) concentrations (Knutti and Hegerl 2008; Lenton et al. 2019).

Yet even if all sources of uncertainty were removed, decisionmakers must still grapple with ethical judgements on how society values climate impacts experienced by individuals across the income spectrum. This is particularly salient for the very poor, who bear little responsibility for historic greenhouse gas emissions but are highly vulnerable to a changing and variable climate (Dennig et al. 2015; Adler et al. 2017; Budolfson et al. 2019). The long timescales associated with climate change add further ethical complexity, raising intertemporal equity concerns about trade-offs between generations and the distribution of policy costs and benefits over time (Schelling 1995; Arrow et al. 1996). Some approaches attempt to infer society's ethical preferences from income tax schedules, foreign aid, and observed market behaviors (Arrow et al. 1996). However, such prescriptive methods are silent on whether the current economy and global distribution of resources represents an ethically defensible basis for climate policy design (Fleurbaey et al. 2019).

One of the most prominent equity concerns in policy analysis involves the value society places on human health and wellbeing. Around the world, ambient air pollution is estimated to cause more than four million deaths per year (World Health Organization 2018). Targeting climate policies towards co-emitting sources of CO₂ and air pollutant emissions could therefore produce important co-benefits through improvements in local health outcomes. However, such policies could also yield substantial co-harms because the aerosols formed from air pollutants produce

beneficial net cooling (G. Myhre et al. 2013). In Chapter 2, I quantify how optimal global CO₂ mitigation pathways change when accounting for the full trade-off between mitigation costs, climate damages, and the climate and localized health consequences of changes in air pollutant co-emissions. I modify a cost-benefit integrated assessment model (IAM) to account for (i) endogenous, income dependent regional aerosol emissions levels, (ii) the changes in radiative forcing due to changes in aerosol emissions, and (iii) the amount of life years gained due to reductions in exposure to ambient air pollution. I then use a utilitarian social welfare function (Fleurbaey et al. 2019) to value the additional life years gained from co-reductions in local air pollution. This framework allows me to analyze how different ethical preferences on inequality affect optimal climate policy decision making. I also create a fully co-optimizing model that allows regions to separately set air pollution and CO₂ mitigation levels and carry out a suite of sensitivity tests to explore the robustness of my results.

In Chapter 3, I pursue a more formal approach for model uncertainty quantification that goes beyond Chapter 2's individual sensitivity tests. The federal government uses the social cost of methane (SC-CH₄) to analyze proposed regulatory actions that will cause changes in methane emission levels (Interagency Working Group on Social Cost of Greenhouse Gases 2016). Despite the policy relevance of the SC-CH₄, few estimates for it exist (Marten et al. 2015) and current estimation approaches neglect key sources of structural and parametric uncertainty in the climate models used. I provide updated, probabilistic SC-CH₄ estimates by pairing IAMs with multiple Bayesian calibrated climate, carbon, and methane cycle models that have been validated through probabilistic hindcast tests. Importantly, my results account for uncertainty in the parameter distribution tails that strongly influence projection uncertainties (Roe and Baker 2007; Urban and Keller 2010; Wagner and Weitzman 2018; Vega-Westhoff et al. 2020) as well as the recent 25% upward revision to methane radiative forcing (Etminan et al. 2016). I also provide the first probabilistic equity-weighted SC-CH₄ estimates. Equity-weighting formalizes the concept that an additional dollar of consumption improves the wellbeing of a poor person more than a rich person (Anthoff, Hepburn, and Tol 2009) and is a primary method for incorporating distributional equity concerns into climate policy analysis.

In Chapter 4, I extend on my SC-CH₄ work to examine how upper tail social cost of carbon (SCC) estimates vary with different climate model calibration approaches and treatments of parametric uncertainties. The U.S. federal government uses the 95th percentile SCC estimate in regulatory cost-benefit analyses to allow for low-probability, high-impact climate change tail risks (Interagency Working Group on the Social Cost of Carbon 2016). In this analysis, I first test the simple climate models found in the three leading IAMs used to calculate official U.S. SCC estimates. I use the default version of these models as well as updated versions that have been calibrated in a Bayesian framework (i) to make probabilistic out of sample hindcast tests of global surface temperature anomalies and changes in sea level, (ii) to test each models temperature response to cumulative carbon emissions, and (iii) to carry out impulse response behavior tests important for estimating the SCC. I then couple the non-climate components of each IAM to a reduced-complexity Earth system model and show how the SCC's upper tail varies when either neglecting interactions between climate model parameters or adopting the

climate uncertainty framework used for official U.S. SCC estimates (Interagency Working Group on the Social Cost of Carbon 2016). Finally, I show how these different uncertainty frameworks influence equity-weighted upper tail SCC values for the United States, Latin America, and sub-Saharan Africa.

CHAPTER 2

The Impact of Human Health Co-Benefits on Evaluations of Global Climate Policy¹

The health co-benefits of CO₂ mitigation can provide a strong incentive for climate policy through reductions in air pollutant emissions that occur when targeting shared sources. However, reducing air pollutant emissions may also have an important co-harm, as the aerosols they form produce net cooling overall. Nevertheless, aerosol impacts have not been fully incorporated into cost-benefit modeling that estimates how much the world should optimally mitigate. Here we find that when both co-benefits and co-harms are taken fully into account, optimal climate policy results in immediate net benefits globally, overturning previous findings from cost-benefit models that omit these effects. The global health benefits from climate policy could reach trillions of dollars annually, but will importantly depend on the air quality policies that nations adopt independently of climate change. Depending on how society values better health, economically optimal levels of mitigation may be consistent with a target of 2 °C or lower.

2.1 Introduction

Climate policies targeting CO₂ may also reduce air pollutant emissions—and the aerosols they produce—as the two share emission sources. Prior studies on the topic have quantified the associated health co-benefits of pre-defined greenhouse gas reduction scenarios (Shindell et al. 2011; 2012; West et al. 2013; Rao et al. 2016; Markandya et al. 2018; Shindell et al. 2018), or estimated the economic impacts from reducing specific pollutants (Bollen et al. 2009; Ikefuji, Magnus, and Sakamoto 2014; Shindell 2015), but these types of impacts have not been fully incorporated into cost-benefit modeling that estimates how much the world should optimally mitigate. We move this literature forward by developing a comprehensive cost-benefit integrated assessment model based on William Nordhaus' Regionalized Integrated Climate Economy (RICE) model, where the new developments allow the model to weigh both the health co-benefits and the climate co-harms of aerosol co-reductions (co-harms exist because aerosols produce net cooling overall (G. Myhre et al. 2013)); the latter in particular has been a largely neglected aspect of the co-benefits discussion. (We use the term co-harm to refer to the net climate harm of aerosol reductions, recognizing that reducing some species of emissions individually may produce different effects than the net of all species together; for example,

¹ This chapter was originally published as:

Scovronick, N.*, Budolfson, M.*, Dennig, F.*, Errickson, F.*, Fleurbaey, M., Peng, W., Socolow, R. H., Spears, D., Wagner, F*. 2019. The impact of human health co-benefits on evaluations of global climate policy. *Nature Communications*. <https://doi.org/10.1038/s41467-019-09499-x>
*authors contributed equally.

The main content of the published paper has been placed in its entirety in the main body of the dissertation.
The supporting information has been placed in its entirety in Appendix A of the dissertation.

although reducing black carbon produces a climate benefit, that effect may be outweighed by the climate harm of reductions in other species.).

These modeling developments, which account for the key air pollutant emissions and their individual properties, provide new capability to investigate fundamental policy questions that have not been answered by existing studies (Clarke et al. 2014). This includes determining: (1) the optimal climate policy across time and how it is affected by independent air quality control, (2) whether climate policy produces immediate net benefits, or if there are intergenerational trade-offs, and (3) if specified climate targets are justifiable on cost-benefit grounds.

To answer these questions, we modify the RICE optimization model (Nordhaus 2010) to include an empirically calibrated, regionally differentiated feedback mechanism whereby reducing CO₂ also reduces regional air pollutant emissions from co-emitting sources. We then quantify and monetize the impact on both health and radiative forcing throughout the world and compute the resulting optimal climate policy. We thus modify the standard trade-off between CO₂ mitigation costs and climate damages with a more complete analysis that simultaneously weighs mitigation costs, climate damages from CO₂, and the health and climate consequences of changes in air pollutant co-emissions. The resulting model estimates optimal climate policy after jointly considering all these factors. As a robustness check, we also modify the widely used FUND (Climate Framework for Uncertainty, Negotiation and Distribution) model (Waldhoff et al. 2014) to include the same mechanisms.

We find that when both co-benefits and co-harms are taken fully into account, optimal climate policy results in immediate net benefits globally, which overturns previous findings from cost-benefit models that omit these effects. The global health benefits from climate policy could reach trillions of dollars annually, but their magnitude will importantly depend on the air quality policies that nations adopt independently of climate change. Depending on how society values better health, we show that economically optimal levels of mitigation may be consistent with a target of 2 °C or lower.

2.2 Methods

The RICE model

The RICE model was first developed in 1996 to analyze the trade-offs between investing in climate mitigation, which incurs a cost relatively soon, and climate damages, which incur costs in the more distant future (Nordhaus and Yang 1996; Nordhaus and Boyer 2000b). RICE is the regionalized counterpart to the DICE model, which is one of three leading cost-benefit climate economy models used by researchers and governments for regulatory analysis, including to estimate the social cost of carbon (Interagency Working Group on the Social Cost of Carbon 2016). Here we describe the key aspects of the standard RICE2010 model; for a more extensive description of this open-access model, see Nordhaus (2010) (Also see Nordhaus and Boyer (2000) for a more extensively documented, but earlier version of RICE).

Briefly, RICE is a regionalized global optimization model that includes an economic component and a geophysical component that are linked. RICE divides the world into 12 regions, some of

which are single countries, while others are groups of countries. Each region has a distinct endowment of economic inputs, including capital, labor, and technology, which together produce that region's gross output via a Cobb–Douglas production function. Pre-mitigation carbon emissions are a function of gross output and an exogenously determined, region-specific carbon intensity pathway. These carbon emissions can be reduced (mitigated) at a cost to gross output through control policies, set to equalize the marginal abatement cost in all regions. Local mitigation cost is borne by each region, and there are no inter-regional transfers. Any remaining (post-mitigation) carbon emissions are incorporated into the climate module where they influence global temperature and, ultimately, the future economy through climate-related damages. Future climate change affects regions differently, with poorer regions generally more vulnerable to climate damages. Damage estimates increase quadratically with a change in the global surface temperature and, like mitigation costs, are incurred directly as the loss of a proportion of gross output. Gross output minus the loss of mitigation costs and climate damages is what we refer to hereafter as GDP.

The model's optimization balances mitigation costs, which lower consumption at the time of mitigation, against climate damages which lower consumption in the future. (Regional consumption is the fraction of GDP that is not saved; mitigation cost and climate damage affect consumption only via their effect on GDP). The optimal trade-off maximizes the sum of discounted well-being, W , which is a concave function of consumption as follows:

$$W(c_{it}) = \sum_{it} \frac{L_{it}}{(1 + \rho)^t} \frac{c_{it}^{(1-\eta)}}{1 - \eta}, \quad (2.1)$$

where L is population, c per capita consumption, ρ the rate of pure time preference, and η the consumption elasticity of marginal utility (inequality aversion). The subscripts i and t are the region and time indices, respectively. The model is solved by maximizing this global objective function. As a result, any factor affecting consumption, such as health impacts or climate damages, can be included in the model's optimization framework. Unless otherwise specified, we maintain RICE's default parameter values for time preference and inequality aversion of 1.5% and 1.5, respectively.

For this study, all simulations were run using the Mimi model development package in the Julia programming language. This Mimi/Julia version of RICE is fully faithful to the standard Excel version, but has more flexibility (Anthoff, Errickson, and Rennels 2018). We make four changes to this standard version of RICE, in addition to those directly related to the AIR module, which are described below. First, we update the population projections to those of the UN2017 medium variant (United Nations Population Division 2017). This is a newer source of projections and also enables us to use internally consistent projections of deaths and life-years, as described further below. The impact of changing population has been comprehensively explored elsewhere (Scovronick et al. 2017; Budolfson et al. 2018; Gillingham et al. 2018). Second, we update the exogenous radiative forcing terms to the values used in RCP6.0 (Meinshausen et al. 2011a), which is in line with the latest versions of DICE (Nordhaus and Sztorc 2013; Nordhaus 2017), which is the global (single-region) variant of RICE. These estimates are newer than those found in RICE2010 and are available in disaggregated form, thus allowing us to remove and endogenize the individual aerosol term for use in the AIR module, while maintaining the other

non-CO₂ forcings (also described in more detail below). Third, we allow CO₂ concentrations and the global temperature to be endogenous in the second time-step (2015–2025) — it is fixed in standard RICE2010 — as the now-endogenized aerosols will produce effects immediately after mitigation. And fourth, we use a modified objective function that avoids Negishi weights, which distort time preferences (Dennig and Emmerling 2017) and the inter-regional trade-off in ways that are opaque and difficult to justify descriptively and normatively (Stanton 2011). We have explained this latter change in more detail in a previous publication (Dennig et al. 2015) and also in a sensitivity presented in Table A.2.

In the standard RICE2010 model, anthropogenic CO₂ is the only endogenous climate forcer. All other sources of radiative forcing, including land-use change, non-CO₂ gases, and aerosols are represented through a single exogenous forcing term that aggregates the individual trajectories of each source. This simplifying assumption is problematic when it comes to aerosols. Mitigation actions affecting CO₂ have the potential to strongly influence emissions of the air pollutants that produce aerosols, as the two types of emissions share many sources (Shindell et al. 2012; West et al. 2013). Therefore, if carbon emissions are reduced, aerosols will tend to decrease simultaneously. A change in aerosols implies a change in radiative forcing as well as a change in ambient particulate air pollution (Shindell et al. 2012). Climate change and air pollution both affect well-being, and capturing the impact of these pathways was the motivation for developing the AIR module, which we now explain.

Overview of the AIR model

In this section, we provide a general overview of how we developed the AIR module, with a technical description—including all equations—in the sections that follow.

Broadly speaking, our approach consists of five steps. First, we estimate the baseline (before carbon mitigation) emissions of five air pollutant species (primary PM_{2.5}, oxides of nitrogen, sulfur dioxide, organic carbon, and black carbon). Emissions are estimated for each region-time pair with income-dependent emission intensity projections (emissions per unit GDP) based on the Greenhouse gas and Air pollution Interactions and Synergies (GAINS) model and specifically the ECLIPSE emission scenarios (Stohl et al. 2015). Our central case assumes air pollutant emissions in the coming decades follow the ECLIPSEV5a baseline scenario, which includes current and planned air quality legislation but no climate policy. In sensitivity analyses we alter this assumption, allowing for faster or slower independent air quality cleanup, including a case where we simultaneously cooptimize both air quality and climate policy. This co-optimization introduces policy levers for end-of-pipe technologies that act on individual air pollutants. These levers require associated cost curves that are also drawn from the GAINS model.

In the second step, we determine the change in air pollutant emissions that would result from a change in CO₂ emissions using information from the SSP project (Rao et al. 2017). This provides an estimate of co-reductions based on empirically realistic projections about the regionally differentiated interaction between future climate and air pollution policies, and is consistent with theoretical results from economics that show that co-benefit effects could be different in scenarios with different properties (Fullerton and Karney 2018). Emission

information in the SSPs is estimated from bottom-up integrated assessment modeling that includes regionally differentiated, spatially explicit representations of energy production and structure and, like in RICE, assumes that CO₂ policy occurs through a single global carbon price.

Third, we link changes in air pollutant emissions to changes in estimated average human exposure to PM_{2.5} by applying the source-receptor matrix (SRM) from the TM5-FASST air quality model (Van Dingenen et al. 2018). The TM5-FASST SRM was computed from simulations of the full TM5 chemical transport model for 56 source regions, which were aggregated to approximate the RICE regions (Van Dingenen et al. 2018). Once exposure is estimated, it is possible to calculate the number of life-years gained attributable to (reduced) air pollution by combining an exposure--response function with projections of future mortality and life expectancy (Pruss-Ustun et al. 2003), which we took from the UN World Population Prospects. We applied a (log)linear exposure-response function for mortality from all causes in adults based on a meta-analysis published in a recent World Health Organization report (Forestiere, Kan, and Cohen 2014). We selected this approach for consistency with the UN projections—which only estimate mortality from all causes—and because recent epidemiological analyses indicate that the strong effects of air pollution occur at exposure (concentration) levels up to and including those most relevant to our study and that they likely affect a wide range of outcomes (Yin et al. 2017; Burnett et al. 2018).

Fourth, we allow the change in air pollutant emissions to influence the global temperature using aerosol forcing coefficients derived from the MAGICC climate model (Meinshausen, Raper, and Wigley 2011). The coefficients incorporate both the direct and indirect effects represented in MAGICC, with the latter including those related to albedo and cloud responses. Aerosol forcing is added to the forcing from the other greenhouse gases in RICE's climate module to produce estimates of future climate change.

And fifth, for each region we monetize and then value the aerosol impacts. The average per capita health benefits are added to per capita consumption, whereas the climate effects — monetized using RICE's standard climate damage function — subtract from consumption (Nordhaus 2010). Consumption is transformed into well-being by a concave function in the optimization via RICE's objective/social welfare function (Eq. 2.1). Once the impacts have been valued in this way, they then enter the optimization.




Figure 2.1: Diagram of the AIR module. Flow chart illustrating how the AIR module (rectangles) links with the RICE model (gray circle) to estimate emissions of air pollutants and their impacts. The

model/method underlying each step in AIR is shown within parentheses, along with the relevant equations.

Estimating baseline emissions

Air pollutant emissions from natural sources and open burning remain exogenous in our framework, following the trajectory based on RCP6.0 (Meinshausen et al. 2011a). All other anthropogenic air pollutant emissions are endogenous, as follows.

The level of baseline (pre-mitigation) anthropogenic emissions of each aerosol precursor, E^0 , is a function of an emission intensity factor (emissions per unit GDP), e , and the GDP, Y :

$$E_{itp}^0 \left(Y_{it}, L_{it}, e_{itp} \left(\frac{Y_{it}}{L_{it}} \right) \right) = e_{itp} \left(\frac{Y_{it}}{L_{it}} \right) \cdot Y_{it} . \quad (2.2)$$

The emission intensity is region (i), time (t), and pollutant (p) specific and characteristically depends on the per capita income level of the region, defined above as the GDP, Y , divided by the population, L . GDP is endogenously determined in RICE while population is exogenous, taking the medium variant estimate of the 2017 version of the UN Population Prospects through 2100 (United Nations Population Division 2017), and remaining constant thereafter.

We estimated emission intensities (and emissions) for five air pollutants: sulfur dioxide (SO₂), primary fine particulate matter (PM_{2.5}), oxides of nitrogen (NO_x), organic carbon (OC), and black carbon (BC). Emission intensities up to the year 2050 were derived from region-specific projections used in the GAINS model and specifically the ECLIPSEV5a scenario that reflects existing national and regional air pollution policies, but excludes decarbonization from climate mitigation (Stohl et al. 2015).

We fit the following functional form to the scenario data to extrapolate beyond 2050:

$$e_{itp} \left(\frac{Y_{it}}{L_{it}}, \chi \right) = \varphi_{1,ip} \cdot \left[\exp \left(-\Omega_{1,ip} \cdot I_{it} \left(\frac{Y_{it}}{L_{it}}, \chi \right) \right) + \varphi_{2,ip} \cdot \exp \left(-\Omega_{2,ip} \cdot I_{it} \left(\frac{Y_{it}}{L_{it}}, \chi \right) \right) \right], \quad (2.3)$$

where the Ω s and φ s are fitting parameters and χ is the Kuznets-relevant income, as described below. Resulting emission intensities are displayed graphically in Fig. A.4.

This functional form implies that emission intensities decrease with rising per capita income I , as observed in the projected emission scenario. This can be interpreted as a particular version of the environmental Kuznets curve, estimated with the ECLIPSE data. We write our Kuznets-relevant income $I(\chi)$ as:

$$I_{it} \left(\frac{Y_{it}}{L_{it}}, \chi \right) = \frac{Y_{i(t=2005)}}{L_{i(t=2005)}} + \chi \cdot \left(\frac{Y_{it}}{L_{it}} - \frac{Y_{i(t=2005)}}{L_{i(t=2005)}} \right). \quad (2.4)$$

For $\chi = 1$, the Kuznets-relevant income equals the true (modeled) per capita income in the region-time pair, producing our default best-guess emission intensity factors. Changing $I(\chi)$ allows us to explore assumptions of more or less stringent autonomous air quality policies: we can speed up ($\chi > 1$) or slow down ($\chi < 1$) the decrease in emission intensities over time accordingly.

Fig. A.5 shows the baseline (pre-mitigation) level of emissions by region and time under the default of $\chi = 1$, where the Kuznets-relevant income equals the true (modeled) income.

Relationship of CO₂ mitigation to air pollutant emissions

The above method projects the level of air pollutant emissions assuming no CO₂ mitigation. Reducing emissions of CO₂ will typically also reduce the emissions of air pollutants, as they often stem from the same sources. In our framework, the percentage reduction in CO₂ relative to the business-as-usual (without mitigation) scenario is called the decarbonization fraction (control rate), μ_{it} , and it is associated through parameter κ_{it} with a reduction in pollutant p relative to its baseline (pre-mitigation) level:

$$\frac{\Delta E_{itp}(\mu_{it})}{E_{itp}^0} = \kappa_{itp} \frac{\Delta E_{it,CO_2}(\mu_{it})}{E_{it,CO_2}^0} = \kappa_{ip} \cdot \mu_{it} , \quad (2.5)$$

or,

$$\Delta E_{itp}(\mu_{it}, Y_{it}, L_{it}) = \kappa_{ip} \cdot \mu_{it} \cdot E_{itp}^0 = \kappa_{ip} \cdot \mu_{it} \cdot e_{itp} \left(\frac{Y_{it}}{L_{it}} \right) \cdot Y_{it} . \quad (2.6)$$

The parameter κ_{ip} describes the effectiveness of CO₂ mitigation in co-reducing emissions of pollutant p . The κ_{ip} parameter was estimated from the SSPs, which are a set of five storylines designed to analyze trade-offs between climate change and socioeconomic factors (Rao et al. 2017). Each of the five SSPs contains multiple sub-scenarios that differ only by their level of decarbonization; all socioeconomic factors remain constant. Therefore, each pairwise comparison of these sub-scenarios includes an implicit estimate of κ_{itp} at each time period for each region and pollutant. The emission information in the SSPs is estimated from bottom-up integrated assessment modeling that includes regionally differentiated, spatially explicit representations of energy production and structure, and like RICE, assumes that CO₂ policy occurs through a single global carbon price.

We fit a simple linear regression line through the multiple estimates of κ for each SSP, constrained to begin at the origin, and take the slope of that line as our estimate (Fig. A.6). We derive five estimates of κ (one for each SSP) for each region-pollutant pair, using the middle-of-the-road SSP2 as our standard case, with the alternative SSPs tested in sensitivity analyses. The total post-mitigation level of emissions is therefore:

$$E_{itp}(\mu_{it}) = E_{itp}^0 - \Delta E_{itp} = E_{itp}^0 \cdot (1 - \kappa_{ip} \cdot \mu_{it}) . \quad (2.7)$$

We acknowledge that there are cases where the data derived from the SSP database appear to exhibit a non-linear rather than linear relationship between CO₂ reduction and air pollutant reduction (Fig. A.6). Our goal was to find a relationship that on average, and in particular at full mitigation, provides a reasonable approximation to the implied air pollution reduction. Having said this, we are aware that our assumption of linearity in the relation may affect our overall results: a higher co-reduction at low CO₂ mitigation offers incentives for further reduction than in the linear case, and vice versa. Thus, the shape of the mitigation profile could be affected, though the effect is likely to be small if, as we assume, the carbon prices in the RICE regions are coupled to each other.

From air pollutant emissions to health impacts

The health co-benefits in RICE + AIR are calculated directly from the change in ambient population-weighted concentrations of PM_{2.5} attributable to CO₂ mitigation. In the next paragraph, we describe how we estimate this change in concentration. Further below we describe how we keep track of the absolute level of pre- and post-mitigation PM_{2.5} concentrations, which are used in the health impact calculations only to ensure that no health benefits accrue at exposures below given threshold values.

The change in ambient concentrations of PM_{2.5} attributable to CO₂ mitigation is a function of the change in aerosol precursor emissions—in this case primary PM_{2.5}, NO_x, and SO₂—as well as other factors such as meteorological conditions. To estimate this relationship, we extracted the SRM from the freely available TM5-FASST global atmospheric SRM (Van Dingenen et al. 2018). For each pollutant, the SRM provides an estimate of the change in population-weighted PM_{2.5} concentrations (hereafter referred to as exposure) given a unit change in emissions. The SRM from TM5-FASST was computed for 56 source regions from simulations with the full TM5 chemical transport model (Van Dingenen et al. 2018). Using an SRM is a practical alternative to running full atmospheric chemistry transport model simulations, which is infeasible in our optimization context. The TM5-FASST model is described in detail elsewhere, and has been used similarly in other projects (Rao et al. 2016; Van Dingenen et al. 2018).

The TM5-FASST interface allows the 56 source regions to be aggregated into larger regions that approximate the RICE regions. Due to the size of the RICE regions, we assume that the change in exposure to PM_{2.5}, ΔC , in region i depends only on the change in emissions within that same region (i') and that the change is estimated via the SRM (SR) that encodes atmospheric transport. We also assume that the SRM does not change over time and that changes in the three precursor pollutants are additive:

$$\Delta C_{it}(\mu_{it}, Y_{it}, L_{it}) = \sum_{i'p} SR_{ii'p} \cdot \Delta E_{i'tp}(\mu_{i't}, Y_{i't}, L_{i't}) = \sum_p SR_{iip} \cdot \Delta E_{itp}(\mu_{it}, Y_{it}, L_{it}). \quad (2.8)$$

As mentioned above, we keep track of the absolute level of exposure, a variable that is used only to ensure that no health benefits accrue at exposures below a given threshold level (described further below). The absolute exposure levels in the pre-mitigation case (where $\mu = 0$) are calculated by translating the change in emissions relative to 2005 into a change in exposure using the SRM, and then subtracting it from the 2005 exposure:

$$C_{it}^0(\mu_{it} = 0) = \max \left(C_{i,t=2005} - \sum_{i'p} SR_{ii'p} \cdot (E_{i',t=2005,p} - E_{i'tp}^0), 0 \right). \quad (2.9)$$

Here the max function ensures PM_{2.5} concentrations do not drop below zero. Emissions and exposure in 2005 were taken from the EDGAR (Emission Database for Global Atmospheric Research) emission database (European Commission Joint Research Centre 2016) and Brauer et al. (2016), respectively, and then aggregated into the RICE regions.

Mitigation of CO₂ ($\mu > 0$) reduces exposure further:

$$C_{it}(\mu_{it}) = \max \left(C_{i,t=2005} - \sum_{i'p} SR_{ii'p} \cdot (E_{i',t=2005,p} - (E_{i'tp}^0 - \Delta E_{i'tp})), 0 \right), \quad (2.10)$$

$$\max \left(C_{i,t=2005} - \sum_{i'p} SR_{ii'p} \cdot (E_{i',t=2005,p} - E_{i'tp}(\mu_{it})), 0 \right). \quad (2.11)$$

We define the health co-benefit as the avoided premature mortality resulting from reductions in PM_{2.5} exposure attributable to CO₂ mitigation (ΔC from Eq. 2.8). This benefit can be quantified through the attributable fraction (AF) (Pruss-Ustun et al. 2003):

$$AF_{it} = \frac{RR_{it} - 1}{RR_{it}}, \quad (2.12)$$

where the relative risk, RR , for each region and each time period is a function of the change in exposure and a health impact function (β) that links a unit change in exposure to a change in the risk of adult (≥ 30) mortality from all causes:

$$RR_{it} = \exp(\beta \cdot \Delta C_{it}(\mu_{it}, Y_{it}, L_{it})). \quad (2.13)$$

We assume a log-linear relationship between PM_{2.5} exposure and all-cause mortality with a relative risk of 1.066 (95% confidence interval (CI) = 1.040, 1.093) for each 10 $\mu\text{g}/\text{m}^3$ change in exposure, based on a meta-analysis published by the World Health Organization (Forestiere, Kan, and Cohen 2014). However, we note that many recent assessments of ambient air pollution have used the cause-specific integrated exposure-response (IER) functions to estimate mortality impacts (Cohen et al. 2017). Here we focus on all-cause mortality for two reasons. First, we have recently shown that population size/growth strongly affects estimates of optimal climate policy, including for reasons unrelated to human health (Scovronick et al. 2017; Budolfson et al. 2018). Therefore, we use the most recent long-term (to 2100) population projections provided by the UN Population Division, which does not publish corresponding estimates of cause-specific mortality. Second, important recent studies indicate that the IER functions may underestimate excess mortality (Yin et al. 2017; Burnett et al. 2018), and suggest that mortality risks at the

exposure levels seen in our regional analyses may fall within the 95% CI of the WHO estimate presented above and tested in sensitivity analyses (Yin et al. 2017).

In all analyses, we assume that components of PM_{2.5} are equally toxic (Forestiere, Kan, and Cohen 2014), that health benefits accrue in the same 10-year time-steps as the improvement in air quality, and that population and mortality remains constant after 2100. We also confine the health impacts to mortality from PM_{2.5} exposure, though note that there is also concern about the health effects of exposure to smaller particles (World Health Organization Regional Office for Europe 2013).

Multiplying the AF by the total number of life-years lost from all causes, Θ , yields the life years gained from the reduction in air pollution:

$$LY_{it}(\mu_{it}, Y_{it}, L_{it}) = AF_{it} \cdot \Theta_{it} \quad (2.14)$$

$$= \frac{RR_{it} - 1}{RR_{it}} \cdot \Theta_{it} \quad (2.15)$$

$$= \left(1 - \exp(-\beta \cdot \Delta C_{it}(\mu_{it}, Y_{it}, L_{it}))\right) \cdot \Theta_{it} \cdot \quad (2.16)$$

Since $\beta \cdot \Delta C_{it}(\mu_{it}, Y_{it}, L_{it}) \ll 1$ we can write as an approximation:

$$LY_{it}(\mu_{it}, Y_{it}, L_{it}) = \beta \cdot \Delta C_{it}(\mu_{it}, Y_{it}, L_{it}) \cdot \Theta_{it} \cdot \quad (2.17)$$

Theta (Θ) values are estimated from the UN data by multiplying the total deaths by the remaining life expectancy at the age of death. As UN life expectancy data is by exact age, reported at 5-year intervals, whereas mortality data is for 5-year age groups, remaining life expectancy for each 5-year age group was taken as the average of the group's bounding ages. For example, remaining life expectancy for all (averted) deaths in the 30–34 age group would be the average of the remaining life expectancy for a 30 year old and a 35 year old.

Health benefits in a given region can accrue until absolute exposure to PM_{2.5} converges to some minimum level, representing either a point below which no additional health impacts occur (a threshold) or a theoretical minimum level where residual PM_{2.5} consists only of natural sources. We followed recent global studies (e.g. Apte et al. (2017)) and chose a lower threshold/theoretical minimum of 5.8 $\mu\text{g}/\text{m}^3$. However, we acknowledge that there is a growing consensus that there may not be a safe level of PM_{2.5} below which no adverse health effects occur (World Health Organization Regional Office for Europe 2013; Di et al. 2017) and therefore ran sensitivities down to 1 $\mu\text{g}/\text{m}^3$. If a reduction in a given time period brings a region's exposure below the threshold, the health co-benefit is calculated from the increment between the unmitigated level and the threshold:

$$LY_{it}^*(\mu_{it}, Y_{it}, L_{it}) = \beta \cdot \Theta_i \cdot \max \left[\min \left(\Delta C_{it}(\mu_{it}, Y_{it}, L_{it}), (C_{it}^0(\mu_{it}, Y_{it}, L_{it}) - \tau) \right), 0 \right], \quad (2.18)$$

where τ is the value of the threshold level.

Radiative forcing and the temperature effects from aerosols

Some air pollutants are climate forcers: BC is a warming agent while SO₂, NO_x, and OC act to cool the atmosphere (G. Myhre et al. 2013). The net global forcing in each time period attributable to aerosols is taken as the sum of the individual contributions:

$$RF_t^{aer} = \sum_p \sum_i r_{ipt} \cdot E_{itp}(\mu_{it}, Y_{it}, L_{it}), \quad (2.19)$$

where r_{ipt} is a region-specific coefficient that relates the regional change in emissions of pollutant p to the change in average global forcing.

We used the MAGICC6 climate model (Meinshausen, Raper, and Wigley 2011) to derive the coefficients by determining the impact on forcing from a pulse change in emissions in each time period in each region:

$$r_{ipt} = \left[\frac{\partial RF_t^{aer}}{\partial E_{itp}} \right]_{MAGICC6}. \quad (2.20)$$

For this we ran MAGICC6 with a pre-defined representative concentration pathway (RCP) scenario and then again after having reduced the emissions of one pollutant in one region by a marginal amount. We repeated this procedure for each region, pollutant, time-step, and, finally, each RCP. We thus derived a reduced-form surface-response representation of aerosol forcing in MAGICC6. We separated the effects of the different pollutants to allow them to be controlled independently, as they are in reality (also see Note A.1). In this experiment, we observed a time dependence that accounts for changes in atmospheric dynamics as emissions accumulate:

$$r_{ipt} = \mu_{1,ip} \cdot t + \mu_{0,ip}. \quad (2.21)$$

The resulting coefficients incorporate both the direct and indirect forcing effects represented in MAGICC6, with the latter including those related to albedo and cloud responses. The time-dependence is independent of the initial conditions. We noted that r depends on the magnitude of the pulse change, but the effect is relatively small. Table A.3 reports values of r_{ipt} and Fig. A.7 shows the radiative forcing over time for three scenarios. Here the reader will note that we have used two separate models for estimating the health effects (TM5-FASST) and climate effects (MAGICC6) of air pollutant emissions, a decoupling that may introduce uncertainty.

Nevertheless, MAGICC and TM5 have distinct strengths that we harness accordingly; MAGICC takes account of the full aerosol load of the whole atmosphere, top to bottom, and at the level of hemispheres, while TM5 tells us about the concentration at ground level, where people live and breathe, in principle at much higher resolution. Both models work with the same global emission inventories.

Aerosol forcing affects global mean atmospheric temperature just as CO₂ forcing does, so that the added atmospheric temperature flow is equal to:

$$\Delta T_t^{atm} = \xi (RF_t^{CO_2} + RF_t^{aer}), \quad (2.22)$$

where $RF_t^{CO_2}$ is CO₂ forcing and ξ is the decadal speed of adjustment for atmospheric temperature (equal to 0.208).

Aerosol feedbacks on the economy

As described, aerosol impacts occur from a change in radiative forcing and from a change in air quality. Changes in radiative forcing are transferred to RICE's climate module where they influence the global average surface temperature, which is the basis of the monetized climate damage estimates (Nordhaus and Boyer 2000b).

Changes in air quality are monetized as the health co-benefit, B , by multiplying the number of life years gained by the value of a life-year (VOLY). We follow the same approach taken in early versions of RICE's climate damage function where a VOLY is assumed to equal 2 years of regional per capita consumption (Nordhaus and Boyer 2000b):

$$B_{it}(\mu_{it}, c_{it}) = VOLY_{it}(c_{it}^{pre-health}) \cdot LY_{it}^*(\mu_{it}, Y_{it}, L_{it}). \quad (2.23)$$

A VOLY of 2 years of per capita consumption is generally the same order of magnitude as empirical estimates based on willingness-to-pay surveys (Desaigues et al. 2011), but we also test several alternative values in sensitivity analyses. Final (post-health) per capita consumption, c_{it} is calculated as:

$$c_{it} = c_{it}^{pre-health} + \frac{B_{it}}{L_{it}}. \quad (2.24)$$

With monetized aerosol impacts now included in the economic framework, RICE can follow its normal optimization procedure to find the decarbonization pathway that maximizes the objective (Eq. 2.1).

2.3 Main Results

We summarize the results of the new model (hereafter referred to as RICE+AIR, for RICE+Aerosol Impacts and Responses) in terms of the optimal fraction of business-as-usual CO₂ emissions that should be reduced over time. We refer to this quantitative reduction in CO₂ as the optimal decarbonization fraction and express it as a percentage. (What we call the decarbonization fraction is often referred to as the control rate.). The blue line in Fig. 2.2a shows the optimal decarbonization pathway if the cost-benefit analysis only considers the climate impacts of CO₂ and associated aerosol co-reductions. Health impacts are not included. In this reference case, the optimal decarbonization fraction is 24% of business-as-usual emissions in 2030, rising to 35% in 2050 and ultimately reaching full decarbonization by 2130. It is

similar to the optimal trajectory of the standard RICE model, which features exogenous aerosol forcing and excludes health co-benefits (Fig. A.1).

In contrast to the blue line, which considers only the climate impacts of CO₂ and aerosols, the red line also includes the health co-benefits of the aerosol reductions that result from decarbonization. This leads to increased optimal CO₂ mitigation. The difference between the red line and the blue line demonstrates the importance of the health benefits; roughly 45–60% more decarbonization is optimal over the next five decades (and 10–40% thereafter) compared to the reference case that only considers climate consequences. Full mitigation also occurs earlier in time. The additional emission reductions that result from the inclusion of the health co-benefits cumulatively amounts to ~270GtC (Fig. 2.2a). Importantly, all of these results account for the damages from lost cooling attributable to the aerosol co-reductions.




Figure 2.2: Optimal decarbonization and temperature. **a**, Decarbonization over time for the reference case optimal policy that excludes health co-benefits (blue line) and in the full RICE + AIR optimal policy that includes health co-benefits (red line). **b**, Estimated global temperature rise above preindustrial levels that would occur given the decarbonization in **a**. Decarbonization is relative to a business-as-usual scenario without any climate action, and 100% decarbonization signifies zero net carbon emissions.




Figure 2.3: Costs and benefits of mitigation. Decomposition of the change in global consumption relative to the business-as-usual (BAU) scenario under **a** the full RICE +AIR optimal policy, and **b** the reference case optimal policy. Health co-benefits and benefits from avoided CO₂ damages are positive, while mitigation costs and aerosol co-harms (climate damage from the co-reduction of cooling aerosols) are negative. The black solid line displays the global net effect. **a** shows that the net effect on global consumption is immediately positive when health co-benefits are taken into account, in contrast to the reference case (**b**), which is representative of standard cost-benefit models that do not include health co-benefits and thus imply that optimal climate policy has net costs for much of this century. If health co-benefits are added to the reference policy in **b**—by adding the light-red bars displayed in **a**—the global net effect becomes immediately positive, and if health co-benefits were removed from **a**, the net effect would be negative for most of this century.

The additional decarbonization justified by the health gains leads to a peak temperature 0.4 °C lower than the reference case (Fig. 2.2b). The carbon price pathways associated with Fig. 2.2 can be found in Fig. A.2.

Figure 2.3a shows that the optimal climate policy has immediate and continual monetized global net benefits when accounting for health co-benefits. This overturns the findings from standard cost-benefit optimization models, which ignore health co-benefits and thus imply that optimal climate policy has net costs for much of this century (Fig. 2.3b). The result is consistent with prior co-benefits studies that have analyzed specific emission reduction scenarios and reported high benefits relative to mitigation costs (Saari et al. 2015; Markandya et al. 2018), but which generally do not also account for (monetized) climate-related impacts.

The distribution of health co-benefits by region for the full RICE + AIR optimum is displayed in Fig. 2.4a. In line with recent scenario-based studies (Li et al. 2018; Markandya et al. 2018; Shindell et al. 2018), many of the co-benefits accrue in India and China in early periods, attributable to their large populations and high capacity for mitigation-related reductions in PM_{2.5}. China's benefits decline by mid-century due to relatively rapid economic development

and a stabilizing population—which both act to constrain emissions—whereas those in India persist and are the major driver of increased decarbonization relative to the reference case (see below for a sensitivity test which corroborates India’s importance). Towards the end of the century, sub-Saharan Africa replaces China as the second-largest beneficiary, as air pollution remains problematic due to lagging economic development accompanied by the world’s largest population. Other regions also stand to benefit, including less populous regions that show important benefits per capita and/or per Gross Domestic Product (GDP) (Fig. 2.4b,d).

In the results presented above, the averted premature mortality from aerosol co-reductions produces annual monetized benefits in the hundreds of billions of dollars over the next few decades, rising to several trillion annually at the end of the century (Fig. 2.4c). To derive these numbers, we multiply, for each region-time pair, the total life-years gained by 2 years of per capita consumption. This approach produces different life-year monetizations for each region, which leads to the slight change in composition of monetized benefits (compare Fig. 2.4a, c). However, this does not imply that we assign life-years in poorer regions less value in the objective function, because the optimization accounts for the diminishing marginal utility of consumption through a concave relationship between wealth and well-being, as described below and in the Methods section (Eq. 2.1). Below we also show the sensitivity of our results to alternative valuations of health benefits.




Figure 2.4: Health benefits of carbon mitigation. Life-years gained **a** overall and **b** per 100,000 population by region from the air quality improvements associated with the optimal decarbonization in

RICE +AIR. **c, d** show the resulting monetized benefits in total, and as a percent of GDP, respectively. Note that if a region's PM_{2.5} exposure (concentration) drops below 5.8 µg/m³, health benefits no longer accrue—this threshold assumption is common in other global air quality assessments (Kolstad et al. 2014), and is tested in the sensitivity analyses.

2.4 Independent Air Quality Control

All results presented above assume the level of air quality control that occurs independently of climate policy will proceed approximately as projected in the coming decades, based on current and planned policies. We systematically test the importance of this assumption by implementing an environmental Kuznets-type approach where emission intensities decrease with increasing per capita income (described in detail in the Methods section, and in particular Eqs. 2.2–2.4 and associated discussion). Figure 2.5 reports results when we slow down and speed up this Kuznets-relevant income, allowing it to range between roughly 50% ($\chi = 0.5$) and 150% ($\chi = 1.5$) of the true (modeled) income; lower values imply less stringent air quality control and vice versa. In these model runs the true income is used in all other parts of the model.

The results demonstrate that although assumptions about air quality control do influence optimal mitigation levels, more decarbonization remains optimal compared to the reference case under all of these scenarios. However, the mechanisms driving the additional decarbonization vary. If the (pre-mitigation) air is dirtier than projected ($\chi < 0.5$), extra mitigation occurs primarily to reap the health co-benefits. If stringent air quality control occurs in the future ($\chi > 0.5$), more decarbonization is also optimal compared to the reference case. In part this is to capture remaining health co-benefits, but also for two other reasons. First, the potential damages associated with the substantial loss of the aerosol cooling effect requires additional CO₂ reductions as a counterbalance. And second, when the air is cleaner, each marginal reduction in CO₂ has a greater net benefit because it is not coupled to as much cooling aerosol.

Another way of understanding the role of independent air quality measures on optimal climate policy is by comparing not only the different RICE+AIR cases, as in Fig. 2.4, but also the difference between the RICE+AIR cases and their corresponding reference cases, as reported in Table 2.1. The results confirm that as autonomous emission controls get stronger (represented by higher χ values), the less health co-benefits matter for climate policy.




Figure 2.5: Impact of independent air quality control. **a** Optimal decarbonization and **b** associated global temperature rise given different levels of autonomous air quality control. The reference and RICE + AIR cases from Fig. 2.1, where $\chi = 1$, are displayed as the blue and red solid lines, respectively.

One of the scenarios displayed in Fig. 2.5 is the fully co-optimal case that selects the ideal combination of both air quality and climate policies by introducing an additional policy lever that can act directly on individual air pollutant emissions through end-of-pipe measures, rather than via CO_2 (described in Note A.1). When this lever is added, large reductions in air pollutant emissions occur when and where associated abatement costs are low, the reductions lead to relatively large decreases in exposure, and/or they produce few climate damages. This first-best policy where climate and air pollution policies are co-optimized leads to rapid air quality improvements but still increased decarbonization relative to the reference case, albeit by a lower margin than in the standard $\chi = 1$ case that underlies our main results. Note that we selected $\chi = 1$ as the basis for the main result in order to explore the effect health co-benefits have when regions act consistently with their current and planned air pollution policies (which are suboptimal). As this section demonstrates, the magnitude of the co-benefit effect is importantly influenced by the assumed level of independent air quality control.

Table 2.1: Impact of independent air quality control.

	2030 Reference	2030 RICEAIR	2050 Reference	2050 RICEAIR	Diff. in peak temp. ($^{\circ}\text{C}$) ¹
$\chi=0.5$	20%	39%	27%	57%	0.9
$\chi=0.75$	22%	37%	33%	54%	0.5
$\chi=1$ (main result)	24%	36%	35%	52%	0.4
$\chi=1.25$	24%	35%	36%	51%	0.4
$\chi=1.5$	25%	35%	37%	50%	0.3

Optimal decarbonization rates in RICE + AIR versus the corresponding reference case for different values of χ , which represent different levels of independent air quality control. $\chi = 1$ represents the assumption underlying the main results, whereas higher (lower) values indicate more (less) stringent air quality control in the future. The difference in peak temperature is the difference between the highest global average temperature in the RICE + AIR optimum versus the reference case given the assumed value of χ .

¹Positive values indicate a lower peak temperature in RICE + AIR. Values are rounded.

2.5 Discounting

The RICE model has a discounted utilitarian objective, meaning that for optimal policy calculations, the objective of the model is to maximize the sum of discounted wellbeing (see Eq. 2.1 of the Methods and associated discussion). The discount rate for consumption is determined via the Ramsey rule, which adds the rate of pure time preference to the product of the elasticity of marginal utility and the economic growth rate. The rate of pure time preference is the rate at which the weight given to future well-being declines over time. The elasticity of marginal utility represents the lesser importance of each additional dollar to well-being as one gets richer. Unless otherwise noted, we assume an elasticity of marginal utility of 1.5, the default in RICE.

In all the results presented so far we have used a rate of pure time preference of 1.5% per year, which is the default value RICE, but close to the upper end of the range highlighted by the IPCC (Intergovernmental Panel on Climate Change) (Kolstad et al. 2014; National Academies of Sciences, Engineering, and Medicine 2017). This choice puts more weight on near-term impacts compared to those occurring further in the future. As a result, aerosol impacts, which occur more immediately than the climate impacts from CO₂, have a relatively outsized importance. Nevertheless, when we implement a much lower rate that corresponds to a near-zero (0.1%) preference for the present over the future, optimal decarbonization remains substantially higher in RICE+AIR compared to the analogous reference case (Fig. 2.6). (Near-zero time preference is often used in the climate economics literature (National Academies of Sciences, Engineering, and Medicine 2017), including in the Stern Review (Stern 2007)).

Conversely, if we increase time preference to 3.5%, which implies a discount rate for consumption of 7%, the difference between RICE and RICE + AIR becomes even more stark. (To calculate the 7% discount rate via the Ramsey rule, we used the economic growth rate in the initial time period. We note that the exact discount rate may change over time due to the effect of climate damages on economic growth; the impact however, is negligible.). We chose the 7% value because it is the highest of the primary discount rates generally used in US federal cost-benefit analyses. (Although most experts consider lower rates to be more appropriate in an intergenerational context such as in the case of climate change-related analyses (Kolstad et al. 2014; National Academies of Sciences, Engineering, and Medicine 2017), the current US administration has signaled a preference for including the higher 7% rate (Executive Office of the President 2017)).

Such high time preference places so much emphasis on the near-term health benefits associated with reducing CO₂ that optimal decarbonization is substantially higher than both the analogous reference case and the reference case with more moderate (1.5%) time preference. These results

indicate that even after strongly discounting the future, a robust climate policy is still warranted. In Fig. A.3 we present further analyses showing results with a 3% discount rate, which is the lower value generally used for US federal cost-benefit analyses.




Figure 2.6: The influence of pure time preference. Optimal decarbonization with low (0.1%), medium (1.5%), and high (3.5%) rates of pure time preference for the reference case (that excludes health co-benefits) and in the full RICE +AIR case (that includes health co-benefits). All runs have a consumption elasticity of marginal utility of 1.5.

2.6 Monetizing and Valuing Health Benefits

The results above assume that one life-year gained equals 2 years of per capita consumption in dollar terms. Two years of per capita consumption corresponds to the approach to monetizing health impacts used in key studies that form the basis of the RICE climate damage function (Tol 2009; Nordhaus 2010), and is similar to survey-based life-year assessments (Desaigues et al. 2011). However, this approach yields much lower monetary benefits than using the value of a statistical life (VSL), which is also widely adopted in the literature. If we assume that each adult death attributable to PM_{2.5} results in approximately 10 years of life lost (Committee on the Medical Effects of Air Pollutants 2010), we would use roughly 8–16 years of per capita consumption per life-year gained, instead of two; thus, in what follows we use 8–16 years of per capita consumption as one possible approximation of a VSL, in addition to a more direct VSL-based approach described below.

The sensitivity of our results to alternative life-year monetizations is reported in Table 2.2. The monetized global health benefits in our main results discussed above would be roughly 4–11 times higher if we used VSL-like monetizations in the optimization, and the associated optimal level of decarbonization would likely be consistent with keeping the maximum global temperature rise to 2 °C. Two degrees is a target specified in the Paris Agreement and widely considered as necessary to avoid dangerous climate change, but one that has generally not been warranted according to previous cost-benefit assessments (using similar discounting parameters)

that omit health co-benefits (Nordhaus 2010). Combining the low (0.1%) rate of time preference with a VSL-based approach may justify a target as low as 1.8 °C.

Table 2.2 reports the sensitivity of the results to differences in how life-years are monetized in dollar terms. In RICE’s cost-benefit approach, an important second step occurs when the monetized health benefits are valued in well-being terms via the objective function (see Eq. 1.1 in the Methods section), which gives the model the aim of maximizing the (discounted) sum of global well-being through time, as is standard in optimal policy modeling. A key feature of the objective function is diminishing marginal utility of consumption, which captures the core concept that an additional dollar generates more well-being when given to a poorer person than to a richer person. The result is that while life-years are assigned a lower dollar amount in poorer countries in absolute terms, they are actually assigned greater value in well-being/utility terms. We discuss this issue in more detail in Note A.2 and show in Table A.1 that the strong effect of adding health co-benefits persists whether life years are valued more highly in poorer regions, less highly, or exactly the same as in wealthier regions; however, the magnitude of the effect changes.

Table 2.2: Monetizing life-years/lives.

	2030	2050	Peak temperature ^a (°C)
Reference (no co-benefits)	24%	35%	3.0
1 year of per capita income	30%	44%	2.9
2 years of per capita income (main result)	36%	52%	2.6
4 years of per capita income	46%	66%	2.3
8 years of per capita income	60%	85%	2.1
16 years of per capita income	78%	96%	1.9
VSL ^b	57%	85%	2.1
VSL with low discounting ^c	79%	100%	1.8

Optimal decarbonization rates in 2030 and 2050, as well as the maximum global temperature rise given different life-year monetizations for health co-benefits. (The last two rows use an alternative method that monetizes lives using the VSL).

^a Above preindustrial levels. Values are rounded.

^b Value of a statistical life. Equation: $VSL = 9,000,000 \times (GDP_{pc}/54,000)^{1.4}$ as in Robinson et al. (2018).

^c In this run, time preference = 0.1%, compared to 1.5% in all other runs.

2.7 Additional Sensitivities in RICE + AIR

Table 2.3 reports results for several other sensitivities, presenting the percent increase in optimal decarbonization rates in the RICE + AIR case compared to the corresponding reference case. The table is organized as follows, with variables in parentheses referring to the relevant parameter in the model equations reported in the Methods.

First (Test 1), we explore different air pollutant co-reduction levels (κ), with the range representing the high and low values after applying alternative estimates from the other four Shared Socioeconomic Pathways (SSPs). Note that the SSPs are not ordered in terms of their co-reduction potentials, but instead reflect different possible futures across multiple socioeconomic dimensions. In the second test (Test 2), we substitute the TM5-FAst Scenario Screening Tool

(TM5-FASST) source-receptor matrix (SRM) for the SRM based on simulations of the European Monitoring and Evaluation Program atmospheric chemistry and transport model. We confined this sensitivity to Asia because it was the only region analyzed as part of a recent project at the International Institute of Applied Systems Analysis. However, the four Asian regions account for ~85% of all life-years gained globally over the next century in the main analyses (Fig. 2.4), and therefore largely drive the findings. In Test 3 and Test 4 we assume that the relative risk for all-cause mortality (β) was the lower and upper bound, respectively, of the confidence interval in Forestiere et al. (2014) rather than the central estimate. Test 5 lowers the PM2.5 threshold below which there are no adverse health effects (τ) to 1 $\mu\text{g}/\text{m}^3$ instead of 5.8 $\mu\text{g}/\text{m}^3$. In Test 6 we assume there is no benefit from reducing air pollution at levels above 50 $\mu\text{g}/\text{m}^3$. We include this counterfactual boundary case for two reasons. The first is to investigate the maximum possible concavity in risk functions at high levels of exposure; in reality however, recent empirical studies (Yin et al. 2017; Burnett et al. 2018) indicate that the marginal effect of air pollution remains positive at levels well above 50 $\mu\text{g}/\text{m}^3$. The second reason is to provide additional emphasis on the importance of India and China, as this test effectively eliminates the impact of both of those nations (as well as the Middle East/North Africa region) over the next several decades and thus explores the model’s sensitivity to their exclusion. Test 7 assumes that climate damages are twice the standard values in RICE, while the final test (Test 8) uses the Finite Amplitude Impulse Response (FAIR) climate model version 1.0 (Millar et al. 2017) as an alternative to RICE’s native climate model (Note A.3 describes how we integrate the FAIR model).

Table 2.3: Additional sensitivity analyses.

	2030	2050	Difference in peak temperature ($^{\circ}\text{C}$) ^a
Main result	52%	49%	0.4
Test 1	7 to 49%	10 to 50%	0.1 to 0.4
Test 2	51%	48%	0.4
Test 3	33%	30%	0.2
Test 4	70%	65%	0.5
Test 5	52%	49%	0.4
Test 6	3%	13%	<0.1
Test 7	30%	29%	0.3
Test 8	82%	76%	0.7

Percent increase in optimal decarbonization rates in RICE + AIR versus the corresponding reference case given different modeling assumptions. The difference in peak temperature is the difference in the highest global average temperature.

^a The difference in peak temperature between the reference case and the RICE + AIR case within each test: positive values indicate a lower peak temperature in the RICE + AIR case. Values are rounded.

2.8 Exploring Model Uncertainty by Linking AIR to the FUND Model

Like RICE/DICE (the Dynamic Integrated Climate Economy model (DICE) is RICE’s global counterpart), the FUND model is another one of the three leading climate economy models used by the US Interagency Working Group to estimate the social cost of carbon (Interagency

Working Group on the Social Cost of Carbon 2016). FUND has different world regions, a different economic framework, a different climate model, and a different specification of damages when compared to RICE, and thus provides an important opportunity to explore model uncertainty. Comparing Fig. 2.2 with Fig. 2.7 demonstrates that results for FUND + AIR and RICE+AIR are qualitatively similar, despite the well-known differences in the structure and policy recommendations of the two models (Interagency Working Group on the Social Cost of Carbon 2016; Adler et al. 2017); adding aerosol impacts leads to dramatically increased optimal levels of decarbonization. (Note A.4 describes how we link FUND to AIR.)




Figure 2.7: FUND + AIR results. Optimal decarbonization rates over time for the FUND reference case (that excludes health benefits) and the FUND+ AIR case (that includes health benefits).

2.9 Discussion and Conclusion

We have developed a new modeling framework for analyzing the costs and benefits of the co-reductions in air pollutant emissions that result from CO₂ policy. We find that these impacts have a critical role in determining optimal decarbonization rates, as the potential health co-benefits that result from improved air quality are large, occur quickly enough to be economically important, outweigh the near-term co-harms from lost cooling, and are concentrated in developing regions. This remains true even with high discount rates and relatively conservative valuations of improved health. However, as the size of the health co-benefits will be partially determined by future air quality policies, decision makers should jointly plan both types of interventions. Depending on how society values health, it may be economically optimal to limit temperature rise to 2 °C or lower, thus corroborating the climate targets from the Paris Agreement. Overall, optimal mitigation results in immediate net benefits globally.

Our findings should be interpreted in light of several factors. First, our sensitivity analyses identified key variables that meaningfully affect the optimal level of decarbonization in RICE + AIR. These include the shape and magnitude of the exposure-response functions that relate PM_{2.5} exposure to mortality, the assumed relationship between CO₂ emissions and air pollutant emissions, the discount rate, and the valuation of the health benefits. The first two factors are

largely amenable to empirical inquiry, whereas the latter two depend partially on ethical judgements; all have a strong bearing on how much society should mitigate, and when.

Second, we follow standard convention in assuming a uniform global carbon price in our optimal policy calculations. Economists prefer this assumption because a uniform price minimizes the cost to the global economy of any particular level of emissions reductions, and is thus a necessary part of the first-best approach to climate policy. However, another necessary part of first-best policy is a non-climate equity policy involving redistribution to address the large economic inequalities that exist throughout the world; in the absence of such an equity-focused policy, a uniform global carbon price would ignore equity considerations, and arguably impose an unjustifiably heavy burden on developing countries. An important feature of our results relevant to this issue is that health co-benefits provide the most incentive for additional decarbonization in lower and middle-income regions. An interesting extension of our research would explore how co-benefits affect mitigation policies in a range of different burden-sharing regimes and in second-best climate policy calculations (Chichilnisky and Heal 1994; Anthoff 2011).

Third, we follow common practice in the co-benefits literature in assuming that there is a background appetite for air pollution policy that is independent of climate policy, such that the same background levels of investment will be made in air quality regardless of CO₂ reductions (and so policies will ratchet up insofar as greater CO₂ reductions make the antecedent level of air quality mitigation less expensive); this is the basic assumption behind the co-reductions we estimate in our main results. As reported above in Fig. 2.5, we also co-optimize air quality and climate policies, which is another scenario (with different properties) in which there would be co-benefits from CO₂ reductions. A limitation is that other scenarios are also possible, such as CO₂ reductions in a context involving a fixed cap and trade regime for air pollutants where the cap remains fixed at the same level in conjunction with new CO₂ reductions; in this scenario, it is possible that there would not be many co-benefits from CO₂ reductions (Fullerton and Karney 2018).

Fourth, we were not able to fully assess uncertainties around key aerosol processes. For example, the aerosol indirect effect may be the single largest source of uncertainty in radiative forcing assessments (Boucher et al. 2013; G. Myhre et al. 2013). New evidence continues to accrue about key factors that contribute to the formation of new particles from precursor gases, and there are likely to be complex feedback mechanisms that occur in response to future temperature changes (Fuzzi et al. 2015; Silva et al. 2016). Some of these processes are not yet adequately represented in even the most sophisticated Earth system models, let alone reduced-form versions. A related concern is that we have used separate models to estimate the health effects (TM5-FASST) and climate effects (Model for the Assessment of Greenhouse gas Induced Climate Change 6 (MAGICC6)) of air pollutant emissions, a decoupling that may also introduce uncertainty. Nevertheless, MAGICC and TM5 have distinct strengths that we harness accordingly; MAGICC takes account of the full aerosol load of the whole atmosphere, top to bottom, and at the level of hemispheres, while TM5 tells us about the concentration at ground level, where people live and breathe, in principle at much higher resolution. Both models work with the same global emission inventories.

Fifth, we have only explored the effect of co-reducing air pollutant emissions on PM_{2.5}-related mortality. Other co-benefits may occur, which would likely push further in the direction of increased decarbonization; these include morbidity impacts from PM_{2.5} — which are generally minor compared to those from mortality — as well as other more indirect impacts such as potential increases in crop yields, improvements in visibility, and health impacts from changing exposure to tropospheric ozone (Shindell 2015; Silva et al. 2016; Chang et al. 2017; Shindell et al. 2018). In addition, cost-benefit models also miss some other impacts of reduced climate change, including the effects of methane on ozone, the effects of ocean acidification, and others (Scovronick, Vasquez, et al. 2019).

Due to the uncertainties inherent in our modeling framework, we expect the accuracy of quantitative estimates of the co-harms and co-benefits of optimal climate policy to sharpen over time, particularly as our understanding of atmospheric science progresses. Nevertheless, the novel modeling approach described here offers important new insights into how much we should mitigate and over what time period, and the sensitivity tests above indicate that we should not expect the qualitative story told by our results to change in light of improved empirical estimates. Our methods also enable investigation of other key questions beyond the scope of this study, including how the inclusion of health co-benefits influences optimal climate policy under different burden-sharing regimes and different worldviews about how much to prioritize the poor, future people, and citizens of other countries.

CHAPTER 3

Improving the Climate Model Calibration Reduces the Social Cost of Methane²

The social cost of methane (SC-CH₄) measures the economic welfare loss caused by emitting one ton of methane into the atmosphere. It plays a crucial role in designing climate policies that target methane emissions. SC-CH₄ estimates are used to value changes in methane emissions in regulatory cost-benefit analysis and they inform the tax rate that should be imposed in an optimal taxation framework. Regulating methane emissions has long been recognized as a key pillar of an economically efficient climate policy design (Reilly et al. 1999; Tol, Heintz, and Lammers 2003; Weyant, de la Chesnaye, and Blanford 2006). Yet, current SC-CH₄ estimates neglect key scientific findings and observational constraints. Here we estimate the SC-CH₄ incorporating the recent 25% upward revision to methane radiative forcing calculations (Etminan et al. 2016), combined with calibrated reduced-form global climate models and an ensemble of integrated assessment models (IAMs). Our multi-model mean estimate for the SC-CH₄ is \$922/t-CH₄ (\$468-1540/t-CH₄, 5-95% range) under Representative Concentration Pathway (RCP) 8.5, a 23% decrease relative to the central estimate recently used by the U.S. federal government under the Obama administration (Interagency Working Group on Social Cost of Greenhouse Gases 2016). Our 95th percentile SC-CH₄ estimate is 52% lower than the corresponding U.S. figure. Under RCP2.6, our multi-model mean falls to \$699/t-CH₄ (\$358-1140/t-CH₄, 5-95% range). These decreases primarily result from tightened equilibrium climate sensitivity estimates paired with the effect of previously neglected relationships between uncertain climate model parameters. Our SC-CH₄ estimates also exhibit sizeable variation across model combinations. While sampling different methane cycle models causes an IAM's estimates to vary at most by 20%, the expected SC-CH₄ can more than double when switching between IAMs. Extending our results to account for societal concerns about equity produces SC-CH₄ estimates that differ by over an order of magnitude between low- and high-income regions. Our central equity-weighted estimate for the U.S. increases to \$8040/t-CH₄ (\$4430-12,300/t-CH₄, 5-95% range) while our estimate for sub-Saharan Africa decreases to \$130/t-CH₄ (\$72-199/t-CH₄, 95% range).

3.1 Introduction

Economically efficient climate policy requires balancing the present costs of reducing methane emissions against future benefits from avoided climate impacts (Tol, Heintz, and Lammers 2003). The SC-CH₄ helps quantify this trade-off by approximating the net present value of economic damages from emitting one ton of methane into the atmosphere. The few existing SC-

² This chapter is adapted from:

Errickson, F., Collins, W.D., Srikrishnan, V., Keller, K., Anthoff, D. 2020. Improving the climate model calibration reduces the social cost of methane. *Working Paper*.

The main content of the paper and methods descriptions have been placed in their entirety in the main body of the dissertation. The supporting information has been placed in its entirety in Appendix B of the dissertation.

CH₄ estimates (Hope 2005; 2006; Marten and Newbold 2012; Waldhoff et al. 2014; Marten et al. 2015) influence a wide array of climate policy decision-making. The California Air Resources Board relied on the SC-CH₄ for evaluating strategies to meet the greenhouse gas emission targets mandated by the California Global Warming Solutions Act of 2006 (California Air Resources Board 2017). At the federal level, the U.S. government uses the SC-CH₄ in cost-benefit analyses and recently estimated that methane emission reductions from proposed oil and natural gas regulations could produce over \$1 billion in climate benefits (United States Environmental Protection Agency: Office of Air and Radiation & Office of Air Quality Planning and Standards 2016) (All comparisons in this paper with official U.S. estimates use the SC-CH₄ figures that were developed and used by the Obama administration⁵. The Trump administration has since withdrawn these estimates and is no longer using them (Executive Office of the President 2017)).

The SC-CH₄'s policy relevance creates a need for SC-CH₄ estimates that carefully reflect our current understanding of the climate system and the corresponding decision-relevant uncertainties. While past SC-CH₄ estimates (Hope 2005; 2006; Marten and Newbold 2012; Waldhoff et al. 2014; Marten et al. 2015) provided important initial insights into the societal harm caused by methane emissions, several major shortcomings make them an inadequate basis for further research and policy design. Specifically, past SC-CH₄ values rely on underestimates of methane's radiative forcing (Etminan et al. 2016) and are silent on the effect important parametric uncertainties and their potential interactions have on climate damage projections (Table B.1). Producing sound SC-CH₄ estimates also requires climate models that enable a careful sampling of the parameter space and are sophisticated enough to replicate key behaviors simulated in comprehensive Earth System models. The climate models found in the cost-benefit integrated assessment models (IAMs) previously used for estimating the SC-CH₄ do not meet this standard, contain well-known flaws (Harmsen et al. 2015; National Academies of Sciences, Engineering, and Medicine 2017), and have not been rigorously calibrated to observations. They may therefore produce biased SC-CH₄ estimates inconsistent with the climate record.

Here, we use an ensemble of IAMs to provide improved, probabilistic SC-CH₄ estimates that account for previously neglected climate uncertainties and the recent 25% upward revision to methane radiative forcing (Etminan et al. 2016). We create four simple, mechanistically motivated models of the coupled carbon, methane, and climate systems by pairing the Simple Non-Linear Earth System (SNEASY) model (Urban and Keller 2010) to the original methane cycle components from: (i) the Finite Amplitude Impulse Response model (FAIR) (Smith et al. 2018), (ii) the Climate Framework for Uncertainty, Negotiation, and Distribution IAM (FUND) (Anthoff and Tol 2014), (iii) the Hector climate model (Hartin et al. 2015), and (iv) the Model for the Assessment of Greenhouse-gas Induced Climate Change (MAGICC) (Meinshausen, Raper, and Wigley 2011). These models are hereafter referred to as *S-FAIR*, *S-FUND*, *S-Hector*, and *S-MAGICC*. Each model utilizes updated radiative forcing equations that account for new insights into methane's shortwave absorption (Etminan et al. 2016) that previous studies ignored. We constrain the models to observations from 1850-2017 using a Bayesian framework. We then sample from the joint posterior parameter distribution, including the tails, which strongly influence projection uncertainties (Ruckert et al. 2017) (see Methods). While we consider a large number of uncertain parameters that may influence the SC-CH₄, this list is by no means exhaustive (Table B.2). We further explore the SC-CH₄'s sensitivity to structural model

uncertainties by coupling each calibrated climate model to the non-climate components of FUND and the Dynamic Integrated model of Climate and the Economy (DICE) (Nordhaus and Sztorc 2013), two leading cost-benefit IAMs previously used by the U.S. government to estimate the SC-CH₄ (we refer to *FUND* as the IAM used for an SC-CH₄ estimate, whereas *S-FUND* refers to the underlying simple climate model). For each climate model, we sample a parameter set from its joint posterior distribution and calculate the SC-CH₄ for a one-ton methane pulse in 2020 under RCP8.5. We repeat this process 100,000 times per climate model-IAM combination to produce 800,000 unique SC-CH₄ estimates.

3.2 Methods

Model coupling

We improve the representation of climate system dynamics in DICE (Nordhaus and Sztorc 2013) and FUND (Anthoff and Tol 2014) by pairing each IAM to new models of the coupled carbon, methane, and global climate systems. We model the dynamics of the carbon cycle and climate with the Simple Non-Linear Earth System Model (SNEASY). SNEASY couples a simple representation of the climate system with a non-linear carbon cycle and has been thoroughly described elsewhere (Urban and Keller 2010). Since SNEASY does not contain an endogenous representation of the methane cycle, we couple it with the methane cycle components from four different models. We describe this procedure in greater detail below. Briefly, SNEASY includes ocean and terrestrial carbon cycle models that respond to temperature changes through a feedback with the climate module (Ricciuto, Davis, and Keller 2008; Urban and Keller 2010). The ocean carbon cycle consists of a four-layer diffusion model, with the top layer representing the atmosphere and mixed ocean surface. The terrestrial carbon cycle accounts for the leafy vegetation, living wood, detritus, and soil carbon sinks.

The climate component is based on the Diffusion Ocean Energy balance CLIMate model (DOECLIM), a globally aggregated energy balance model that couples a zero-dimensional atmosphere to a one-dimensional diffusion model of the ocean (Kriegler 2005; Tanaka and Kriegler 2007). While DOECLIM distinguishes between land masses and the ocean, the low dimension is, of course a drastic approximation to the complex spatial patterns of climate change. The global mean surface temperature represents an area-weighted average of land and sea surface temperature anomalies resulting from an induced radiative forcing.

In the original version of SNEASY, radiative forcing from sources other than changing atmospheric carbon dioxide concentrations enter the model exogenously. We create four new versions of SNEASY that endogenously estimate changes in methane's atmospheric concentration and the resulting global surface temperature response by coupling SNEASY to the methane cycle components from: FAIR (Smith et al. 2018), FUND (Anthoff and Tol 2014), Hector (Hartin et al. 2015), and MAGICC (Meinshausen, Raper, and Wigley 2011) models (accessed from the Global Change Assessment Model source code; <https://github.com/JGCRI/gcam-core>). We further update each model to use improved radiative forcing equations for carbon dioxide and methane that importantly account for the recent upward revision to direct methane radiative forcing estimates (Etminan et al. 2016). For methane's indirect effects on stratospheric water vapor and tropospheric ozone, we use each model's

original radiative forcing equations. We also follow each model's original approach for quantifying the additional carbon dioxide produced by fossil fuel-based methane oxidation; S-FAIR and S-MAGICC account for this oxidation effect while S-FUND and S-Hector do not.

We make two additional model modifications. First, we convert MAGICC, which takes methane concentrations as an input before the year 2000 and then switches to methane emission inputs, to use methane emissions for all time periods. This allows us to use methane concentration data from before the year 2000 as an observational constraint. Second, the original FAIR model uses a fitted, time-varying natural methane emission scenario to ensure the model output matches historic atmospheric methane concentration values (Smith et al. 2018). After 2005, the model switches to a constant natural methane emission rate for all subsequent years. We replace FAIR's fitted natural methane emissions scenario with a constant natural emission rate (following the approach adopted by Hector and MAGICC) to avoid artificially inflating the model's atmospheric methane concentration hindcast skill. For each model, we treat annual natural methane emission rates as an uncertain parameter. The calibration jointly samples values for these emission rates alongside the other uncertain methane cycle and climate parameters (Table B.2) to produce model output consistent with the set of observational constraints.

Calibration and forcing data

We use annual observational time series data of the climate system to calibrate each model over the period 1850-2017. Each dataset provides measurement error estimates that enter our statistical calibration framework (described below). The data include atmospheric carbon dioxide concentrations from Mauna Loa (Keeling et al. 1976; Thoning and Tans 1989) and Law Dome ice cores (Etheridge et al. 1996), chlorofluorocarbon based estimates of oceanic carbon fluxes (McNeil et al. 2003), the HadCRUT4 global temperature data set (Morice et al. 2012), and measurements of 0-3000 meter depth ocean heat uptake (Gouretski and Koltermann 2007). Although SNEASY has a 4000-meter ocean, we assume minimal heat exchange occurs below 3000 meters during the calibration time horizon (Urban et al. 2014). The calibration also utilizes atmospheric methane concentrations from globally averaged marine surface data (Dlugokencky 1994) and Law Dome ice cores (Etheridge et al. 1998). Following past work (Etheridge et al. 1998), we calculate the globally averaged atmospheric methane concentration from the ice core data as the Law Dome measurement plus 37% of the annual interpolar methane concentration difference.

We run the models using greenhouse gas emission and radiative forcing values from RCP8.5 (Meinshausen 2010; Meinshausen et al. 2011a), a no-policy scenario that closely tracks current carbon dioxide emission levels (Le Quéré et al. 2018; Smith and Myers 2018), but helps highlight differences in methane cycle model behavior by providing large forcing and emission signals in future years. In a sensitivity test, we repeat our analysis using the RCP2.6 scenario. Each model exogenously accounts for radiative forcing effects from solar and volcanic activity, direct and indirect effects of aerosols, land use albedo changes, stratospheric ozone, black carbon on snow, and anthropogenic greenhouse gases other than carbon dioxide or methane. The greenhouse gas forcing values include fluorinated gases regulated by the Kyoto protocol in addition to all gases controlled under the Montreal Protocol.

Unlike the other models in this study, SNEASY-F (SNEASY + the FUND methane cycle) approximates methane's indirect radiative forcing from stratospheric water vapor and tropospheric ozone as 40% of methane's direct radiative forcing (Anthoff and Tol 2014; Interagency Working Group on Social Cost of Greenhouse Gases 2016). However, SNEASY-F does not endogenously calculate total tropospheric ozone radiative forcing. Simply running SNEASY-F with or without the RCP8.5 tropospheric ozone radiative forcing values would therefore produce biased parameter estimates and model projections. To address this issue, we use the FUND methane cycle model to derive an exogenous, RCP8.5 consistent radiative forcing scenario for tropospheric ozone's non-methane component.

- Step (1): Calculate the tropospheric methane lifetime that minimizes the root-mean-square error (RMSE) between modeled atmospheric methane concentrations and RCP8.5 concentrations over the calibration period 1850-2017.
- Step (2): Use the RMSE minimizing methane lifetime and RCP8.5 methane emission values to project atmospheric methane concentrations out to 2300.
- Step (3): FUND calculates methane's indirect radiative forcing due to tropospheric ozone as 25% of methane's direct radiative forcing (Interagency Working Group on Social Cost of Greenhouse Gases 2016; IPCC 2007). Using this factor and the concentration time series from Step (2), calculate FUND's predicted methane contribution to tropospheric ozone radiative forcing.
- Step (4): Subtract FUND's predicted methane contribution from the RCP8.5 total tropospheric ozone radiative forcing time series. This difference serves as an exogenous input to SNEASY-F for the non-methane component of tropospheric ozone radiative forcing.

Bayesian model calibration

To derive statistically sound parameter estimates that provide results consistent with the historic climate record, we calibrate each model using a Bayesian framework. The joint posterior distribution $p(\theta|y)$ for climate model j represents the probability of observing the uncertain parameters θ after accounting for information contained in the climate observations y and comes from Bayes' Theorem:

$$p(\theta_j|y) = \frac{p(\theta_j)L(y|\theta_j)}{p(y)} . \quad (3.1)$$

Here, the prior distribution $p(\theta_j)$ expresses prior knowledge about the uncertain parameters and the likelihood function $L(y|\theta_j)$ provides the likelihood of observing the data given model j 's parameters. The denominator $p(y)$ represents a normalizing constant that often requires numerical approximation, making it difficult to perform Bayesian inference on even relatively simple climate models. We overcome this issue by using the robust adaptive Markov chain Monte Carlo (MCMC) algorithm within our Bayesian framework (Metropolis et al. 1953;

Hastings 1970; Gilks, Richardson, and Spiegelhalter 1996; Vihola 2012). This approach allows us to sample $p(\theta_j|y)$ by calculating Eq. 3.1 up to proportionality, making it a computationally tractable model calibration approach (Smith and Roberts 1993; Gilks, Richardson, and Spiegelhalter 1996; Vihola 2012).

We adopt physically-informed truncated uniform prior distributions for all uncertain model parameters except for the equilibrium climate sensitivity ECS , the vertical diffusivity of heat into the ocean K_v , a multiplicative scaling factor on aerosol radiative forcing α , and the initial global surface temperature anomaly T_0 (Table B.2). For these parameters, we follow previous work and adopt subjective prior distributions (Urban and Keller 2010). The ECS prior uses a truncated Cauchy distribution to account for paleoclimate information not utilized in the calibration while K_v uses a log-normal prior to reflect data derived from biogeochemical tracers (Urban and Keller 2010). We use a triangular distribution for α to approximate IPCC estimates on the total radiative forcing effects of aerosols. Finally, we use a standard normal distribution centered on zero to describe uncertainty about T_0 . The prior distributions for all uncertain parameters are assumed to be independent, with their product giving the joint prior distribution from Eq. 3.1.

We calculate the likelihood function from Eq. 3.1 in terms of the residuals between the observed and modeled climate. Past work shows that assuming too simple of an error structure in the statistical model fitting framework can produce biased parameter estimates and overconfident model projections, particularly for low-probability tail events (Zellner and Tiao 1964; Ricciuto, Davis, and Keller 2008; Ruckert et al. 2017). We address this issue by approximating the residuals using autoregressive process models that also account for time-varying observation errors contained in the data (Ruckert et al. 2017). For time series n of observed data y at time t , model j 's residual ε_{jnt} represents the difference between the model output $f(\theta)_{jnt}$ and observations:

$$f(\theta)_{jnt} - y_{nt} = \varepsilon_{jnt} . \quad (3.2)$$

The data-model residuals in Eq. 3.2 consist of two distinct components:

$$\varepsilon_{jnt} = e_{jnt} + \omega_{nt} , \quad (3.3)$$

where the autoregressive model error e_{jnt} represents the potential residual autocorrelation and ω_{nt} represents time-varying observation error. We assume the observation errors come from a mean-zero normal distribution with a time-varying variance $\sigma_{nt\ obs}^2$ given by measurement error estimates provided with each calibration dataset $\omega_{nt} \sim \mathcal{N}(0, \sigma_{nt\ obs}^2)$.

The global mean surface temperature anomaly and oceanic heat content time series both contain annually consecutive observations. We hence model the corresponding residuals using a stationary first order autoregressive process, or AR(1), model. In this framework, the value of e_{jnt} depends linearly on the previous residual through a first-order autocorrelation coefficient ρ_{jn} that varies across datasets but remains constant over time:

$$e_{jnt} = \rho_{jn} \times e_{jn(t-1)} + \delta_{jnt} . \quad (3.4)$$

Eq. 3.4 also contains an independent and identically distributed (i.i.d.) stochastic error term $\delta_{jnt} \sim i. i. d. \mathcal{N}(0, \sigma_{jn}^2_{AR1})$, with $\sigma_{jn}^2_{AR1}$ the AR(1) innovation variance for data set n . δ_{jnt} captures remaining structural model errors and natural climate variability unable to be resolved by the simplified climate models used in this study. We treat ρ_{jn} and $\sigma_{jn}^2_{AR1}$ as uncertain statistical process parameters and estimate them during model calibration.

The covariance matrix Σ_{jn} for the global mean surface temperature anomaly and oceanic heat content time series represents the sum of the AR(1) process and time-varying observation error variances:

$$\Sigma_{jn} = \frac{\sigma_{jn}^2_{AR1}}{1 - \rho_{jn}^2} \begin{pmatrix} 1 & \rho_{jn} & \rho_{jn}^2 & \dots & \rho_{jn}^{k-1} \\ \rho_{jn} & 1 & \rho_{jn} & \dots & \rho_{jn}^{k-2} \\ \rho_{jn}^2 & \rho_{jn} & 1 & \dots & \rho_{jn}^{k-3} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \rho_{jn}^{k-1} & \rho_{jn}^{k-2} & \rho_{jn}^{k-3} & \dots & 1 \end{pmatrix} + \begin{pmatrix} \sigma_{n1}^2_{obs} & 0 & 0 & \dots & 0 \\ 0 & \sigma_{n2}^2_{obs} & 0 & \dots & 0 \\ 0 & 0 & \sigma_{n3}^2_{obs} & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & \sigma_{nk}^2_{obs} \end{pmatrix}. \quad (3.5)$$

The atmospheric carbon dioxide and methane concentration time series both contain a combination of irregularly spaced ice core and more recent instrumental observations. We model the corresponding residuals with a stationary first-order continuous time autoregressive, or CAR(1), model. A detailed description of this approach can be found in Koen (2005). The CAR(1) model contains a continuous time white noise process with variance $\sigma_{jn}^2_{CAR1}$ and a correlation memory process that decays exponentially with time and is characterized by the term α_{jn} . We treat $\sigma_{jn}^2_{CAR1}$ and α_{jn} as additional uncertain parameters and estimate them during model calibration.

The covariance matrix for the atmospheric carbon dioxide and methane concentration time series then represents the sum of the CAR(1) process and time-varying observation error variances (Koen 2005):

$$\Sigma_{jn} = \frac{\sigma_{jn}^2_{CAR1}}{2\alpha_{jn}} \begin{pmatrix} 1 & e^{-\alpha_{jn}(|t_1-t_2|)} & \dots & e^{-\alpha_{jn}(|t_1-t_k|)} \\ e^{-\alpha_{jn}(|t_2-t_1|)} & 1 & \dots & e^{-\alpha_{jn}(|t_2-t_k|)} \\ \vdots & \vdots & \ddots & \vdots \\ e^{-\alpha_{jn}(|t_n-t_1|)} & e^{-\alpha_{jn}(|t_n-t_2|)} & \dots & 1 \end{pmatrix} + \begin{pmatrix} \sigma_{n1}^2_{obs} & 0 & \dots & 0 \\ 0 & \sigma_{n2}^2_{obs} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_{nk}^2_{obs} \end{pmatrix}, \quad (3.6)$$

where e is the natural exponential function, and $|t_n - t_k|$ describes the time interval between potentially irregularly spaced observations for years n and k . For this analysis, we compare each Law Dome atmospheric carbon dioxide and methane concentration observation with the eight-year average of modeled concentrations, centered on each ice core observation's corresponding year (Urban and Keller 2010).

For the observational datasets mentioned above, we use a likelihood function that assumes the full residual time series $\vec{\epsilon}_{jn}$ comes from a k -dimensional multivariate normal distribution

(Ruckert et al. 2017) characterized by the appropriate covariance matrix $\vec{\epsilon}_{jn} \sim \mathcal{N}_k(0, \Sigma_{jn})$. We write the likelihood function²¹ for dataset n as:

$$L(\vec{y}_n | \theta_j, \Sigma_{jn}) = \left(\frac{1}{\sqrt{2\pi}} \right)^{k_n} |\Sigma_{jn}|^{-\frac{1}{2}} \exp \left[-\frac{1}{2} (\vec{y}_n - \vec{f}(\theta)_{jn})^\top \Sigma_{jn}^{-1} (\vec{y}_n - \vec{f}(\theta)_{jn}) \right]. \quad (3.7)$$

The ocean carbon flux dataset lacks annually consecutive observations. We therefore neglect any potential autocorrelation and apply a likelihood function that assumes normally distributed i.i.d. residuals. We hold the ocean carbon flux standard deviation constant at the reported observational error (McNeil et al. 2003) of 0.4 Pg C/yr.

Under the simplifying assumption that residuals across data sets are independent, the product of the individual likelihoods gives the total likelihood of the observational climate data.

For each version of SNEASY, we set the MCMC target distribution to the numerator of Eq. 3.1 (i.e. the product of the prior and likelihood) and produce a Markov chain that constitutes a representative sample from the joint posterior parameter distribution. We check for convergence to the target distribution using graphical diagnostics and potential scale reduction factor calculations (Gelman et al. 1995). Each chain contains six million parameter samples and we discard the first one million sampled parameters for the burn-in period. We then thin the remaining chain down by selecting a subset of 100,000 parameter samples at equally spaced intervals.

Estimating the social cost of methane (SC-CH4)

For each posterior parameter sample and corresponding IAM combination, we produce a unique SC-CH4 estimate following a four-step process. In the first step, we select a single posterior parameter sample for a given model and calculate two climate trajectories out to 2300; (1) a baseline model run following the RCP8.5 scenario and (2) a pulse run that adds one extra ton of methane emissions in 2020 but is otherwise identical to the baseline run. The resulting climate projections contain superimposed noise to account for process and observation uncertainty (Ruckert et al. 2017). For posterior parameter sample i and climate model j , we simulate this noise term for the n th calibration data set variable (e.g. annual global surface temperature anomalies) by sampling from a mean-zero multivariate distribution characterized by the covariance matrices from Eqs. 3.5 and 3.6. We assume years occurring outside the periods with data coverage have observation measurement errors equivalent to the average of the ten nearest observation measurement errors in time.

In the second step, we pair the uncertain climate projections with an IAM to calculate climate damage estimates. This produces two damage trajectories, a baseline trajectory and one that accounts for the additional impacts caused by the one-ton methane pulse in 2020.

In the third step, we use IAM k to calculate annual marginal climate damages MD_{kjit} as the difference in damages D_{kjit} between the baseline and methane pulse trajectories:

$$MD_{kjit} = D_{kjit(pulse)} - D_{kjit(baseline)} . \quad (3.8)$$

Because the DICE model runs on 5-year timesteps (Nordhaus and Sztorc 2013), we linearly interpolate the DICE damage estimates in Eq. 3.8 to produce annual values.

In the fourth and final step, we calculate the SC-CH4 for a given posterior parameter sample and model combination as the discounted sum of marginal damages across r regions and over time:

$$SC\ CH4_{kji} = \sum_{t=0}^T \sum_{r=1}^R MD_{kjitr} \times \frac{1}{(1 + \Phi)^t} , \quad (3.9)$$

where Φ represents a constant consumption discount rate. We set Φ to 3% in our main results to compare with the central estimates currently used by the U.S. federal government, which are based on a constant 3% rate (Interagency Working Group on Social Cost of Greenhouse Gases 2016). We further align our results with reported U.S. SC-CH4 values by expressing our estimates in 2007 U.S. dollars using conversion factors derived from U.S. Bureau of Labor Statistics inflation data (United States Department of Labor: Bureau of Labor Statistics 2020).

We repeat this entire process for each posterior parameter sample and corresponding climate-IAM model pair. With two IAMs, four climate models, and 100,000 posterior parameter samples per climate model, this yields 800,000 unique SC-CH4 estimates (Fig. 3.2a).

Bayesian model averaging

To address uncertainty in climate model selection, we employ a Bayesian model averaging (BMA) approach that places greater weight on models better able to produce results consistent with past observations (Hoeting et al. 1999; Wong et al. 2018). More formally, the SC-CH4 estimates corresponding to the j th climate model M_j are now weighted based on that model's posterior probability:

$$BMA_j = p(M_j|y) = \frac{p(y|M_j)p(M_j)}{\sum_{q=1}^4 p(y|M_q)p(M_q)} , \quad (3.10)$$

where BMA_j represents the BMA weight for climate model j and we assume each candidate climate model's structure has an equal prior probability ($p(M_j) = p(M_l)$ for all j, l). Making model dependence explicit, the probabilities $p(y|M_j)$ are given by:

$$p(y|M_j) = \int_{\theta_j} p(y|\theta_j, M_j)p(\theta_j)d\theta_j , \quad (3.11)$$

where θ_j corresponds to the calibrated parameters for climate model j . We follow previous work (Meng and Wong 1996; Wong et al. 2018) by approximating Eq. 3.11 using bridge sampling and the posterior results obtained during model calibration with MCMC. The multi-model BMA-weighted SC-CH4 estimate can then be calculated as:

$$SC\ CH4_{BMA} = \sum_{j=1}^4 \overline{SC\ CH4}_j \times BMA_j, \quad (3.12)$$

where $\overline{SC\ CH4}_j$ represents the SC-CH4 averaged across DICE and FUND using climate model j .

Alternative SC-CH4 estimates

We provide four alternative SC-CH4 estimates to explore how our results change under different model structures and treatments of parametric uncertainties. We discard a small subset ($< 0.3\%$) of parameter samples in these sensitivity analyses that produce model errors.

- (1) *Using outdated methane radiative forcing equations.* We recalculate the SC-CH4 for each climate model-IAM pair using the original methane radiative forcing equations found in each methane cycle's parent model (Fig. 3.2b, Fig. B.2). For FAIR, which provides the option to run the model under a variety of forcing assumptions, we select the older radiative forcing equations that do not account for methane's shortwave absorption (Gunnar Myhre 1998; Smith et al. 2018). Each SC-CH4 estimate therefore comes from a climate model that neglects the recent 25% upward revision to methane radiative forcing (Etminan et al. 2016).
- (2) *Removing calibrated climate parameter relationships.* Several posterior relationships between the uncertain parameters emerge during model calibration. These relationships act as an important constraint on model behavior (Fig. 3.3b and c, Fig. B.5), climate projection uncertainty (Fig. 3.3a, Fig. B.6), and the resulting SC-CH4 distributions (Fig. 3.2b, Fig. B.2). For this analysis, we recalculate the SC-CH4 and remove these posterior relationships by assuming the calibrated climate and statistical process parameters come from independent distributions. We then sample from each marginal posterior distribution without replacement to produce new SC-CH4 estimates.
- (3) *Sampling equilibrium climate sensitivity only.* Selecting an appropriate ECS value represents a persistent source of uncertainty for the simplified climate models found in IAMs (Marten 2011; Hof et al. 2012; Harmsen et al. 2015). Federal U.S. SC-CH4 estimates represent this uncertainty by sampling an ECS distribution parameterized to match several broad IPCC probability statements about the ECS (Interagency Working Group on the Social Cost of Carbon 2010; Interagency Working Group on Social Cost of Greenhouse Gases 2016). We replicate this estimation strategy within our framework by fixing each uncertain parameter to its mean posterior value and sampling the ECS distribution used for U.S. SC-CH4 estimates (Roe and Baker 2007; Interagency Working Group on the Social Cost of Carbon 2010) (Fig. B.7a). We calculate a new SC-CH4 value (Fig. 3.2b, Fig. B.2) for each ECS value l sampled from:

$$ECS_l = \frac{1.2}{1 - f}, \quad (3.13)$$

where f comes from a normal distribution $f \sim N(0.62, 0.034)$ truncated between -0.2 and 0.88 to capture uncertainty in climate system feedbacks (Marten and Newbold 2012).

(4) *Wider prior parameter distributions.* Using uniform or truncated prior distributions in the model calibration precludes sampling parameter values beyond their distribution bounds. We carry out a sensitivity analysis using wider prior distributions that allows the MCMC algorithm to explore a much larger parameter space (Table B.2). Based on these wider prior distributions, we recalibrate each of the four climate models to obtain new posterior parameter samples and then produce updated SC-CH4 estimates for each climate model – IAM pair (Fig. B.4a).

Equity-weighting

As discussed above, the SC-CH4 calculated from Eq. 3.9 uses a constant consumption discount rate to facilitate comparison with official U.S. SC-CH4 values (Interagency Working Group on Social Cost of Greenhouse Gases 2016). This approach neglects the concept that an additional dollar of consumption more strongly affects the welfare of a low-consumption individual than a high-consumption individual. To address this concern, we use the FUND model to produce equity-weighted SC-CH4 estimates that also incorporate preferences on risk aversion (Anthoff, Hepburn, and Tol 2009). We produce a new point estimate for each posterior parameter sample to derive equity-weighted SC-CH4 distributions. Under this framework, the equity-weighted SC-CH4 for climate model j and posterior parameter sample i is given by:

$$Equity\ SC\ CH4_{jix} = \sum_{t=0}^T \sum_{r=1}^R \left(\frac{c_{ji0r}}{c_{ji0x}} \right)^{-\eta} \left(\frac{c_{jitr}}{c_{ji0r}} \right)^{-\eta} (1 + \rho)^{-t} MD_{jitr} , \quad (3.14)$$

where η represents the consumption elasticity of marginal utility, ρ is the pure rate of time preference MD_{jitr} represents marginal damages for region r in year t as described in equation (9), and c gives regional per capita consumption. The x subscript identifies the normalization region used in the equity-weight term $\left(\frac{c_{ji0r}}{c_{ji0x}} \right)^{-\eta}$ to normalize the SC-CH4 into a welfare-equivalent loss of consumption for that specific region. The equity-weighted SC-CH4 therefore varies with x and depends heavily on the normalization region's consumption level in the present.

The underlying social welfare function from which Eq. 3.14 derives is a standard utilitarian social welfare function:

$$SWF = \sum_{t=0}^T \sum_r^R P_{tr} U(c_{tr}) \left(\frac{1}{1 + \rho} \right)^t . \quad (3.15)$$

We omitted the climate model and posterior parameter sample index for clarity. We use

$$U(c) = \begin{cases} \log c & \text{iff } \eta = 1 \\ \frac{c^{1-\eta}}{1-\eta} & \text{iff } \eta \neq 1 , \end{cases} \quad (3.16)$$

as the utility function, a standard choice in the literature. The details of the derivation of Eq. 3.16 can be found in Anthoff, Hepburn, and Tol (2009).

For our main results, we set the parameters in Eq. 3.14 to values commonly found in the climate economics literature (Anthoff, Hepburn, and Tol 2009), namely $\eta = 1.0$ and $\beta = 1.0\%$. In a sensitivity test, we also provide results across a range of η values (Fig. 3.5). Note that our results do not account for income inequality or income-dependent vulnerability to climate impacts at the sub-regional level, which could potentially further increase the spread in equity-weighted SC-CH4 estimates between low and high-consumption regions (Dennig et al. 2015).

Uncertainty reporting

The primary objects of interest are the posterior distributions of key model parameters and the predictive distributions for quantities like the SC-CH4, the latter being the decision relevant uncertainties. When we report uncertainties in this chapter, we generally attempt to characterize the shape of these distributions. For example, when we report credible intervals for key results, we are reporting the 5% and 95% percentile of these distributions. For reported multi-model results, we assign each model an equal weighting and average across the individual model estimates. In addition, percent ranges in the main text correspond to the spread of percentage changes in the expected SC-CH4 across all individual climate model-IAM pairings. Percent values are always relative to a model's baseline expected SC-CH4 estimate (RCP8.5 under a constant 3% consumption discount rate). We report our SC-CH4 results to three significant figures.

3.3 Observations Constrain SC-CH4 Estimates

We test our calibrated climate models through probabilistic hindcasts of surface temperature, ocean heat content, ocean carbon flux, as well as atmospheric carbon dioxide and methane concentrations (Fig. 3.1, Fig. B.1). The hindcasts capture the spread of global surface temperature anomaly observations within their 95% predictive credible intervals reasonably well (Fig. 3.1a-d), a particularly relevant test because temperature increases drive DICE and FUND's climate damage projections. The atmospheric methane concentration hindcasts vary somewhat across models (Fig. 3.1e-h), with the simplest model, S-FUND, visibly exhibiting the least skill (Fig. 3.1f).




Figure 3.1: Probabilistic climate hindcast skills vary with model choice and treatment of parametric uncertainties. Hindcasts over the model calibration period (1850–2017) using the S-FAIR (purple), S-FUND (teal), S-Hector (red), and S-MAGICC (orange) climate models. **a–d**, The annual average global surface temperature anomaly relative to the 1861–1880 mean, with yellow circles representing the temperature observations used in the calibration. **e–h**, The annual average atmospheric methane concentration, with different yellow shapes identifying the Law Dome ice core (triangles) and globally averaged marine surface (circles) atmospheric methane concentration observations. In all panels, solid black center lines depict the mean response across 100,000 model runs, with colored regions spanning the 95% predictive credible interval. Gray shading in the temperature hindcasts shows the 95% predictive credible interval for an experiment that samples the ECS distribution used for U.S. SC-CH₄ estimates while fixing other uncertain climate model parameters at their mean posterior values (see Methods).

We use the calibrated climate models to estimate a multi-model mean SC-CH₄ of \$922 (\$468–1540, 5–95% range) (2007 US dollars per metric ton CH₄, \$/t-CH₄) under a constant 3% consumption discount rate (Fig. 3.2a). This represents a 23% decrease relative to the central U.S. estimate under the same discounting framework (Interagency Working Group on Social Cost of Greenhouse Gases 2016), though we note this U.S. value falls within our 95% SC-CH₄ predictive credible interval. While a constant consumption discount rate facilitates a better comparison with official U.S. SC-CH₄ estimates, it fails to account for linkages between economic growth and discounting (National Academies of Sciences, Engineering, and Medicine 2017). In a sensitivity analysis, we recalculate the SC-CH₄ under a standard Ramsey discounting framework (Arrow et al. 2014) and obtain similar estimates (Table B.3). Our results also remain largely unchanged when considering a wider parameter space (Fig. B.4a).

Our SC-CH₄ estimates are rather sensitive to forcing and emission scenario assumptions. Disregarding the recent upward revision to methane radiative forcing calculations (Etminan et al. 2016) decreases our central SC-CH₄ estimate by 16% (Fig. 3.2b, Fig. B.2). Switching from the RCP8.5 to RCP2.6 emissions scenario yields even larger changes in the SC-CH₄. The lower background temperatures projected under RCP2.6 (Fig. B.3a–d) cause an additional methane emission pulse to be less damaging to society and decrease the multi-model mean SC-CH₄ to \$699/t-CH₄ (\$358–1140/t-CH₄, 95% range) (Fig. B.3e). This represents a 24% reduction relative

to our RCP8.5 estimate, an effect comparable in magnitude to changes produced by some of the different discounting assumptions we consider (Table B.3).




Figure 3.2: SC-CH₄ distribution estimates. **a**, Main SC-CH₄ distribution estimates using each calibrated climate model for RCP8.5 under a constant 3% consumption discount rate. Different colors identify the four climate models. White shapes along the x-axis show the estimated mean for FUND (circle) and DICE (diamond) when pooling together estimates from the four climate models. **b**, Alternative SC-CH₄ experiments for the S-MAGICC climate model using a constant 3% consumption discount rate under RCP8.5. Colors identify each SC-CH₄ experiment, with our main specification using Bayesian calibration (“Main SC-CH₄ estimate”, red), outdated radiative forcing equations that disregard methane’s shortwave absorption (“Outdated CH₄ radiative forcing”, orange), neglecting posterior relationships and sampling each parameter independently (“Remove parameter relationships”, green), and sampling the ECS distribution used for U.S. SC-CH₄ estimates while fixing other uncertain climate model parameters at their mean posterior values (“U.S. Climate Sensitivity”, blue). Colored shapes depict the estimated means for the different distributions. In both **a** and **b**, different distribution line types distinguish between the FUND (solid) and DICE (dashed) IAMs.

3.4 Non-Linear Parameter Relationships

As expected from physical reasoning and consistent with previous studies (Forest et al. 2002; Urban and Keller 2010), several strong and non-linear relationships between the model parameters emerge from the calibrations. These relationships prove vital for the climate model’s ability to produce results consistent with the observational record. For example, high equilibrium climate sensitivity (ECS) values are associated with increased aerosol cooling and more rapid ocean heat penetration that counteracts some additional warming that would otherwise occur (Fig. 3.3b, Fig. B.5a-c). Carbon cycle interactions further enrich this relationship. High ECS values are preferentially sampled alongside reductions in the sensitivity of terrestrial carbon pools to increasing temperatures. Strong posterior relationships also arise between uncertain methane cycle parameters. For instance, higher natural emission rates tend to partially offset cases where methane has a shorter atmospheric lifetime (Fig. B.4b). Perhaps surprisingly, we

find little to no association between these uncertain methane cycle parameters and the SC-CH₄ itself (Fig. B.4c). This stands in stark contrast to the ECS and aerosol cooling strength parameters (Fig. 3.3c, Fig. B.5d-f), suggesting the SC-CH₄ responds more strongly to parametric uncertainties that directly influence temperature dynamics rather than methane cycle behaviors.

Neglecting these posterior relationships (see Methods) as done in previous studies (Hope 2005; 2006; Marten and Newbold 2012; Waldhoff et al. 2014; Marten et al. 2015) increases temperature projections (Fig. 3.3a, Fig. B.6) and SC-CH₄ uncertainties (Fig. 3.2b, Fig. B.2). While our expected SC-CH₄ estimates increase by 5-7% in this scenario, the standard deviation of a model's SC-CH₄ estimates increase by 51-55%.




Figure 3.3: Non-linear climate parameter relationships constrain probabilistic temperature projections and SC-CH₄ estimates. a, Modeled annual average global surface temperature anomalies relative to the 1861-1880 mean using the S-MAGICC climate model. Yellow circles represent observations and solid black center line depicts the mean response across 100,000 model runs. Colored

regions outlined by dashed lines span the 95% predictive credible interval. The outer gray colored regions bound the 95% predictive credible intervals for model projections that remove parameter relationships by sampling each marginal posterior distribution independently. The inset figures depict estimated temperature distributions for 2050 and 2100 using the full set of model runs from the baseline (colored) and no parameter relationship (gray) scenarios. **b**, Posterior relationships between four uncertain climate parameters for the S-MAGICC climate model. Different sized circles correspond to different terrestrial carbon pool respiration-temperature sensitivity values (with higher values signaling increasing heterotrophic respiration of carbon dioxide with temperature). Each dot's color scales with the vertical rate of heat diffusion into the ocean. **c**, Posterior relationships between three of the parameters depicted in **b** with the SC-CH₄ estimated using the S-MAGICC climate model for RCP8.5 under a constant 3% consumption discount rate. Different sized diamonds and circles identify respiration temperature sensitivity values for DICE and FUND. Each point's color scales with the aerosol radiative forcing factor (with higher values signaling stronger aerosol cooling). Both **b** and **c** depict 5,000 randomly selected posterior estimates, with loess-smoothed curves (white lines) helping illustrate the relationship between x and y-axis values.

Our lower mean SC-CH₄ estimate arises primarily from different treatments of these parametric climate uncertainties, with the ECS playing an important role. U.S. SC-CH₄ estimates reflect an ECS distribution consistent with several broad statements about likely ECS ranges from the Intergovernmental Panel on Climate Change's Fourth Assessment Report (Interagency Working Group on the Social Cost of Carbon 2010; Interagency Working Group on Social Cost of Greenhouse Gases 2016). Notably, this approach neglects the non-linear parameter relationships described above while also failing to constrain the ECS's deeply uncertain upper distribution tail (Knutti and Hegerl 2008). When we replicate this aspect of the U.S SC-CH₄ estimation strategy (see Methods), the climate models produce upwardly biased temperature hindcasts (Fig. 3.1a-d). This upwards bias is consistent with the expected effects of (i) the higher probability of sampling a high ECS (Fig. B.7a), and (ii) neglecting the posterior parameter relationships. The association between the ECS and SC-CH₄ also strengthens, with linear correlation coefficients equal to or exceeding 0.95 across all model combinations (Fig. B.7b-e). As a result, the expected SC-CH₄ increases by 13-24% and more closely aligns with the official U.S. estimates (Fig. 3.2b, Fig. B.2). The 95th percentile SC-CH₄, used in U.S. federal cost benefit analyses to represent high impact climate scenarios, exhibits considerable sensitivity to the upper ECS distribution tail and increases by 60-78% (Table B.3).

3.5 Model Structure and SC-CH₄ Variation

The results above weight each climate model equally and hence neglect key aspects of model structural uncertainty. Yet differences in hindcast skill suggest assigning equal model weights would be a poor assumption. We employ a Bayesian model averaging (BMA) framework (Wong et al. 2018) that weights our SC-CH₄ estimates based on climate model posterior probabilities (see Methods) and find the resulting BMA weights exhibit considerable variation across the four methane cycle models. More than 80% of the weight is given to S-MAGICC (0.89), followed by S-Hector (0.06), S-FAIR (0.04), and S-FUND (≈ 0.001). However, these BMA weights produce negligible changes to the SC-CH₄ and increase the multi-model mean by only 4%. This occurs, in part, because sampling different methane cycles has only a modest effect on the SC-CH₄'s

expected value and distribution. For instance, switching the climate model from S-FUND (which consistently provides the lowest expected SC-CH₄ estimates) to S-MAGICC (which often provides the highest) increases each IAM's expected SC-CH₄ by roughly 20%. In contrast, switching IAMs from FUND to DICE for a given climate model increases the expected SC-CH₄ by approximately 160%. This suggests non-climate differences between DICE and FUND produce greater variation in the SC-CH₄ compared to the structural uncertainty of the underlying methane cycle models considered here.

The differences separating DICE and FUND become starker when examining the time profile of their climate damage response to a one-ton methane emission pulse. Across multiple constant discount rates, DICE's projected annual marginal discounted damages peak between roughly \$15-45 within a decade of the pulse and then slowly converge towards zero (Fig. 3.4b). In contrast, FUND often produces negative initial discounted damage estimates, projecting that society benefits from additional methane emissions in early periods (Fig. 3.4c). These benefits, which include increased agricultural productivity and avoided heating demand (Diaz and Moore 2017), partially explain FUND's consistently lower SC-CH₄ estimates relative to DICE. However, in both IAMs the fading temperature pulse (Fig. 3.4a) paired with economic discounting pushes damage projections towards zero and leads to decreasing uncertainty over time. This effect becomes stronger under higher discount rates (Fig. 3.4b-c, inset).




Figure 3.4: Discounted climate damage impulse response behavior and uncertainties strongly differ between DICE and FUND. **a**, Temperature impulse response behavior for the S-MAGICC climate model from a one-ton methane pulse in 2020. Solid black center line depicts mean response across 100,000 model runs and outer dashed black lines bound the 95% predictive credible intervals. Orange lines show 1,000 randomly sampled individual model runs. Colored circles identify the corresponding temperature distributions in the inset figure for the years 2030, 2040, 2050, and 2070 estimated from 100,000 model runs. **b**, Discounted climate damage impulse response behavior for DICE paired with the S-MAGICC climate model from a one-ton methane pulse in 2020. In each year, the predicted climate damages are discounted based on constant consumption discount rates of 2.5% (orange), 3% (red), 5% (blue), and 7% (green). For each discount rate, lines outlined in black show the mean model response while colored lines show 250 randomly sampled individual model runs. The inset figure shows discounted damage distributions estimated from 100,000 model runs under 2.5% and 5% constant

discount rates for the years 2030, 2040, 2050, and 2070, with each year distinguished by a different line type. **c**, same as in **b** but depicting results for FUND.

3.6 Equity-Weighting

All of the results, thus far, use a constant consumption discount rate to facilitate comparison with official U.S. SC-CH₄ estimates (Interagency Working Group on Social Cost of Greenhouse Gases 2016). This approach is silent, however, on the question of equity. For example, a given consumption loss reduces a poor person's well-being more than a rich person's. Equity-weighting accounts for this effect by giving more weight to climate damages occurring in lower-consumption regions. This approach is consistent with standard economic theory (Anthoff, Hepburn, and Tol 2009) as well as the call for common but differentiated responsibilities and respective capacities (Page 2008). Equity-weighted social cost of carbon estimates receive considerable research attention and are currently used by the German government (Matthey and Bunger 2019). However, only a single published equity-weighted SC-CH₄ estimate exists to our knowledge and this estimate does not account for the effects of climate model uncertainties (Waldhoff et al. 2014).

Considering equity-weighting offers new perspectives, but also adds complexities. Our equity-weighting framework uses a standard utilitarian social welfare function that assigns equal importance to the well-being of different individuals, but also captures that a dollar is "worth" more to a poor person than a rich person through the normative parameter of *elasticity of marginal utility of consumption* (η , see Methods). In an uncertain setting, η also represents preferences on risk aversion. To illustrate the importance of η for our equity-weighted SC-CH₄ estimates, consider two individuals in a single time period with consumption levels of \$1,000 and \$10,000. When $\eta = 1.0$, a \$1 decrease in consumption produces a welfare loss ten times greater for the poorer individual relative to the richer one. Increasing η to 1.5 causes the poorer individual's welfare loss to be roughly thirty times greater. Using a constant consumption discount rate, as we do throughout this paper to match the U.S. SC-CH₄ estimation approach, is equivalent to setting η to zero and implicitly adopts the ethical stance that \$1 of forgone consumption has the same impact on well-being for rich and poor individuals alike.

In the selected equity-weighting framework, we first aggregate climate damages from an additional ton of methane emissions as the sum of discounted welfare losses across regions and over time. The resulting SC-CH₄ estimate (measured in welfare terms) can then be monetized as a welfare equivalent consumption loss *for a given normalization region* in the present (Anthoff, Hepburn, and Tol 2009). Because socioeconomic inequities exist between regions and additional consumption has less "worth" as one gets richer, the equity-weighted SC-CH₄ varies with the choice of normalization region. Rich regions show higher equity-weighted SC-CH₄ than poor ones. For similar reasons, equity-weighting can yield different regional methane tax rates in an optimal taxation framework if consumption equalizing transfers between regions are either impossible or politically infeasible (Anthoff 2011; Budolfson et al. 2019). A complementary policy framework could also use equity-weighted SC-CH₄ estimates to help identify an initial allocation of methane emissions permits accounting for equity, with permit trading between regions leading to additional economic benefits (Shiell 2003; Budolfson et al. 2019).

We generate equity weighted SC-CH₄ distributions for FUND and report estimates for different choices of normalization region (DICE is a global model and can therefore not be used to compute equity-weighted results, see Methods). As expected, we find substantial variation based on each region's socioeconomic status (Fig. 3.5). Due to its higher consumption levels, the United States' expected equity-weighted SC-CH₄ is \$8040/t-CH₄ (\$4430-12,300/t-CH₄, 95% range) when η equals 1.0, nearly sixty times higher than sub-Saharan Africa's value of \$130/t-CH₄ (\$72-199/t-CH₄, 95% range) (Fig. 3.5c). However, the spread between these estimates increases non-linearly with η (Fig. 3.5a and b). Increasing η to 1.5 as done in the example above increases the United States' SC-CH₄ to \$35,300/t-CH₄ (\$21,000-52,300/t-CH₄, 95% range), now over four hundred times higher than sub-Saharan Africa's \$73/t-CH₄ (\$43-107/t-CH₄, 95% range) estimate.




Figure 3.5: Variations in equity-weighted SC-CH₄ estimates and their uncertainties for different consumption elasticity of marginal utility values. All estimates produced with the S-MAGICC climate model for RCP8.5 using a 1.0% pure rate of time preference. **a**, Equity-weighted SC-CH₄ estimates for five FUND regions as the consumption elasticity of marginal utility increases from 0 to 1.5 (higher values correspond to increasing preferences for equality in consumption). Different colors depict different regions, with the outer edges of each color representing the equity-weighted SC-CH₄ 95% predictive credible interval. Centered lines outlined in black show the expected equity-weighted SC-CH₄. The dashed horizontal gray line bounds all estimates that fall below \$4000/t-CH₄. **b**, Zoomed-in view of the estimates in **a** falling below the \$4000 boundary. **c**, Equity-weighted SC-CH₄ distributions when the consumption elasticity of marginal utility equals 1.0. These distributions correspond to the vertical slice in **a** identified by the red arrow. Colored circles show the estimated means of the different distributions. To improve graph legibility, distributions for the sub-Saharan Africa and China + (China, Hong Kong, North Korea, Macau, and Mongolia) regions have been cropped. The dashed line distribution depicts a

case where the SC-CH₄ estimates do not vary across regions because the consumption elasticity of marginal utility equals 0.

3.7 Discussion and Caveats

The large SC-CH₄ variations caused by equity concerns, different forcing scenarios, and IAM choice suggest other non-climate factors can play a vital role in estimating the SC-CH₄. For instance, climate impact projections strongly depend on IAM damage functions as well as the model's ability to capture society's future climate adaptation efforts. These model elements contain a number of uncertainties (Diaz and Moore 2017; National Academies of Sciences, Engineering, and Medicine 2017) that are not incorporated into our SC-CH₄ estimates. While we account for some structural uncertainty about the climate damage function by using two different IAMs, it is important to note that DICE and FUND do not reflect the full range or capture the highest damage estimates found in the climate impacts literature (Carleton and Hsiang 2016). We also do not consider potential interactions between an IAM's uncertain climate and economic parameters. Relatedly, the U.S. SC-CH₄ incorporates estimates from the PAGE model, while ours do not. This may bias our results downward (Diaz and Moore 2017); however our qualitative findings still hold when restricting the comparison to the U.S. results produced by DICE and FUND alone (Interagency Working Group on Social Cost of Greenhouse Gases 2016). As FUND is generally producing lower damage estimates than DICE, our equity-weighted results may also be biased downwards. More generally, the SC-CH₄ does not account for local mitigation co-benefits (Scovronick, Budolfson, et al. 2019) as well as methane's direct effect on public health and agricultural productivity (Shindell 2015). Both are important additional factors to consider when designing methane emission policies, but are conceptually distinct from the SC-CH₄.

Our results have direct implications for policies targeting methane emissions. Most importantly, they show improved and carefully constrained climate models produce economically meaningful changes to the SC-CH₄ and reduce it relative to official U.S. estimates. This reduction persists despite the 25% upward revision to methane radiative forcing and is particularly evident in estimates from the SC-CH₄'s deeply uncertain but decision-relevant upper tail. Our equity-weighted estimates can yield even larger variations in the SC-CH₄ between low- and high-income regions, but importantly depend on society's tolerance for inequality and risk. Extending our work to explore the joint relationship between climate and socioeconomic uncertainties as well as the complex interactions that arise when accounting for equity represents a promising and policy-relevant avenue for future research.

CHAPTER 4

Enhanced Representation of Climate Model Uncertainties Reduces Upper Tail Social Cost of Carbon Estimates³

The social cost of carbon (SCC) informs how society navigates the trade-off between the upfront costs of reducing greenhouse gas emissions and the benefits of avoiding climate change impacts. The U.S. government uses a central SCC estimate in federal regulatory impact analyses of proposals affecting carbon dioxide (CO₂) emission levels, but also applies the SCC's 95th percentile value to account for the possibility that climate change impacts may end up in the upper tail of the currently estimated distribution (Interagency Working Group on the Social Cost of Carbon 2016). Past estimates of the SCC's upper tail rely, however, on outdated climate models (National Academies of Sciences, Engineering, and Medicine 2017; Dietz et al. 2020) and neglect observational constraints that strongly influence century-scale temperature and sea level projections (Urban and Keller 2010; Flato et al. 2013; Bakker et al. 2017; National Academies of Sciences, Engineering, and Medicine 2017; Vega-Westhoff et al. 2020). Here, we couple the DICE, FUND, and PAGE integrated assessment models (IAMs) with a calibrated, reduced-complexity Earth system model to quantify how addressing these shortcomings impacts the SCC's upper tail. Our multi-model mean SCC estimate is \$49/t-CO₂ (\$26-65/t-CO₂, 5-95% range) for a CO₂ emission pulse in 2020 under Representative Concentration Pathway (RCP) 8.5. The upper tails of these SCC estimates are highly sensitive to parameter correlations that emerge from the model calibrations. Neglecting the correlations, as done in previous SCC estimates, produces physically implausible parameter combinations that increase the 95th percentile by 23%. Under the climate uncertainty framework used for official U.S. SCC estimates (Interagency Working Group on the Social Cost of Carbon 2016), our 95th percentile estimate nearly doubles.

4.1 Introduction

Deciding on how much and when to reduce greenhouse gas emissions depends in part on uncertainties affecting long-term temperature and sea level projections. Although the severity of climate change impacts depends on the total warming that eventually occurs (Field et al. 2014; Carleton and Hsiang 2016), key temperature metrics such as the equilibrium climate sensitivity (ECS) have been persistently uncertain for decades (Knutti and Hegerl 2008; Meehl et al. 2020; Sherwood et al. 2020). For the world's coastal communities, decision makers grapple with end-of-century sea level projections that can differ by more than a meter due to the deeply uncertain behaviors of the Greenland and Antarctic ice sheets (Bakker et al. 2017; Bamber et al. 2019; Oppenheimer et al. 2019). The presence of these uncertainties raises the question of how they should be accounted for in policy metrics such as the social cost of carbon (SCC). The SCC

³ This chapter is adapted from:

Errickson, F., Wong, T.E., Keller, K., Anthoff, D. 2020. Enhanced representation of climate model uncertainties reduces upper tail social cost of carbon estimates. *Working Paper*.

The main content of the paper and methods descriptions have been placed in their entirety in the main body of the dissertation. The supporting information has been placed in its entirety in Appendix C of the dissertation.

helps to inform the decision of how to navigate the trade-off between near-term mitigation costs and avoided climate damages. The United States (U.S.) federal government accounts for low-probability, high-impact tail risks by including the 95th percentile SCC estimate in regulatory cost-benefit analyses (Interagency Working Group on the Social Cost of Carbon 2010). In the Obama administration's analysis of the U.S. Clean Power Plan, switching from the central to the 95th percentile SCC yielded an additional \$60 billion of climate benefits between 2020-2030 (United States Environmental Protection Agency: Office of Air and Radiation & Office of Air Quality Planning and Standards 2015). This example illustrates how sound estimates of the SCC's upper tail can be critical for assessing different climate policy options.

Recent work provides important insights on shortcomings in the simple climate models used for estimating the SCC. Key issues include missing or incomplete carbon cycle-temperature feedbacks and unrealistic climate response timescales to changes in radiative forcing (National Academies of Sciences, Engineering, and Medicine 2017). As a result, policy-relevant discrepancies arise between the impulse response behaviors produced by these simple models and the dynamics observed in more sophisticated Earth system models. For instance, the simple climate models exhibit a multi-decade delay in their temperature response to a large, instantaneous CO₂ emissions pulse that places them out of line with current scientific understanding (Joos et al. 2013; Ricke and Caldeira 2014; National Academies of Sciences, Engineering, and Medicine 2017; Dietz et al. 2020). This delayed behavior biases the SCC downwards because the climate impacts from an additional ton of CO₂ occur further in the future and are discounted more heavily (Dietz et al. 2020; Hänsel et al. 2020).

Additional sources of climate model uncertainty may potentially influence the impulse response behaviors relevant for estimating the SCC. While past model diagnostic experiments have provided important new insights, they generally focus on a single model parameterization and do not explore the full range of uncertain parameters used for estimating the SCC (Tables C.1-3) or their potential interactions (Forest et al. 2002; Urban and Keller 2009; 2010; Errickson et al. 2020). In addition to temperature, cost-benefit integrated assessment models (IAMs) also use sea level projections to estimate damages such as permanent land submergence, more extreme storm surges, and increased coastal adaptation and resiliency costs. Although these damages directly influence the SCC and represent welfare losses for more than 600 million people living along the world's coastlines, surprisingly little attention has been paid to sea level impulse response behaviors. Similar to their temperature counterparts, unrealistic sea level response behaviors could bias SCC estimates. As an example, failing to account for deeply uncertain threshold responses such as the potential rapid disintegration of the West Antarctic ice sheet ignores the risk that society could face larger sea level damages earlier than anticipated (Bamber et al. 2009; DeConto and Pollard 2016; Bakker et al. 2017; Wong, Bakker, and Keller 2017; Frederikse et al. 2020).

Here, we adopt a Bayesian calibration framework to quantify how different treatments of parameter uncertainties affect climate model behavior and the SCC uncertainties. We create modified versions of the simple climate models found in the Dynamic Integrated model of Climate and the Economy (DICE) (Nordhaus and Boyer 2000a; Anthoff 2019b), the Climate Framework for Uncertainty, Negotiation, and Distribution (FUND) (Anthoff 2019a), and the Policy Analysis of the Greenhouse Effect (PAGE) (Hope 2011; Moore et al. 2018) IAMs. The

modifications allow us to run each model over the historic period and calibrate them to past observations. We refer to these models as *DICE-Climate*, *FUND-Climate*, and *PAGE-Climate*. Owing to the well-known flaws in the simple IAM climate models described above (National Academies of Sciences, Engineering, and Medicine 2017; Dietz et al. 2020), we also improve the representation of the climate system using a reduced-complexity Earth System model that couples the Simple Non-Linear Earth System model (SNEASY) (Urban and Keller 2010) to the Building blocks for Relevant Ice and Climate Knowledge model (BRICK) (Wong et al. 2017). SNEASY pairs a one-dimensional diffusive energy balance model to a non-linear carbon cycle. BRICK individually models sea level contributions from glaciers and small ice sheets, land water storage, thermal expansion, and the Greenland and Antarctic ice sheets. We refer to this coupled model as *SNEASY+BRICK*.

We carry out a two-stage model testing strategy (see Methods). First, we calibrate an expanded set of uncertain parameters for each model (Tables C.1-4) to observations spanning the 1850-1970 period. We then use these calibrated parameters to make probabilistic out-of-sample predictions of global mean sea level and the global average surface temperature anomaly from 1971 to 2017. Second, we recalibrate each model to the full set of observations (1850-2017) and carry out a suite of probabilistic temperature and sea level impulse response diagnostics for a 100 GtC emissions pulse occurring in 2020. At each testing stage, we perform sensitivity analyses to analyze how the hindcasts and impulse response behaviors change when we (i) neglect the parameter correlations that emerge from model calibration and (ii) follow the uncertainty framework used for official U.S. SCC estimates (see Methods). To estimate the SCC, we couple SNEASY+BRICK to the non-climate components of DICE, FUND, and PAGE and refer to them as *SB-DICE*, *SB-FUND*, and *SB-PAGE*. We sample 10,000 parameters from SNEASY+BRICK's joint posterior distribution and use them to calculate the SCC for a CO₂ emissions pulse in 2020 under the no-policy RCP8.5 scenario.

4.2 Methods

Model Descriptions

We create modified versions of the simple climate models found in the Dynamic Integrated model of Climate and the Economy (DICE) (Nordhaus and Boyer 2000a; Anthoff 2019b), the Climate Framework for Uncertainty, Negotiation, and Distribution (FUND) (Anthoff 2019a), and the Policy Analysis of the Greenhouse Effect (PAGE) (Hope 2011; Moore et al. 2018) integrated assessment models (IAMs). The modifications allow us to calibrate these simple climate models to past observations, perform probabilistic out-of-sample prediction tests, and carry out a suite of model behavior diagnostics described below. We refer to these models as *DICE-Climate*, *FUND-Climate*, and *PAGE-Climate* throughout the manuscript.

We also couple the Simple Non-Linear Earth System model (SNEASY) (Urban and Keller 2010) to the Building blocks for Relevant Ice and Climate Knowledge model (BRICK) (Wong et al. 2017). We refer to this coupled model as *SNEASY+BRICK*.

SNEASY+BRICK

SNEASY models temperature dynamics based on the Diffusion Ocean Energy balance CLIMate model (DOECLIM), which couples a zero-dimensional atmosphere to a one-dimensional ocean heat diffusion model (Kriegler 2005; Tanaka and Kriegler 2007). SNEASY also contains non-linear terrestrial and ocean carbon cycle models that interact with DOECLIM through a temperature feedback (Ricciuto, Davis, and Keller 2008; Urban and Keller 2010).

BRICK projects changes in global mean sea level as the sum of individually modeled contributions from land water storage, thermal expansion, glaciers and small ice caps, the Greenland ice sheet, and the Antarctic ice sheet (Wong et al. 2017). We make several changes to BRICK before coupling it with SNEASY:

- (i) The original version of BRICK would run each component separately over the full projection time horizon. It would then normalize all results relative to a common period and combine them in one final step. We modify BRICK so all components initialize together and then run in sequential timesteps to facilitate the model's use within IAMs.
- (ii) We incorporate the potential for rapid Antarctic ice sheet disintegration to occur after crossing an uncertain temperature threshold (Wong, Bakker, and Keller 2017) as well as an updated thermal expansion component based on changes in ocean heat content (Vega-Westhoff et al. 2019).
- (iii) For multi-century projections under high forcing scenarios, the Greenland ice sheet model could produce implausibly rapid melt rates. The e-folding timescale of changes in Greenland ice sheet volume τ originally depended on global surface temperature only. Following the approach taken in BRICK's representation of glaciers and small ice caps (Wigley and Raper 2005; Bakker, Applegate, and Keller 2016; Wong et al. 2017), we modify τ 's equation to also depend on the remaining ice sheet volume:

$$\frac{1}{\tau_t} = (\alpha_{GIS} * Temp_t + \beta_{GIS}) * \left(\frac{V_t}{V_0}\right). \quad (4.1)$$

Here, α_{GIS} is the temperature sensitivity of the Greenland ice sheet exponential decay rate, $Temp_t$ represents the global surface temperature anomaly at time t , and β_{GIS} captures the exponential decay rate for a 0 °C global surface temperature anomaly.

The newly added term $\left(\frac{V_t}{V_0}\right)$ represents the ratio between the Greenland ice sheet's remaining and initial volumes. It serves to slow the speed of decay as the ice sheet's volume approaches zero.

DICE, FUND, and PAGE climate models

We first annualize DICE-Climate and PAGE-Climate so they can be calibrated to the full set of historic climate observations described below. In contrast to FUND-Climate, which runs on annual timesteps, DICE-Climate runs for ten year periods and PAGE-Climate has irregular time

steps for the years 2009, 2010, 2020, 2030, 2040, 2050, 2075, 2100, 2150, and 2200. The annualization process required minimal modifications to each model's default structure. For DICE-Climate, we removed any factors of ten that scaled the model output to cover a decade-length period. PAGE-Climate automatically scales its output to account for the variable lengths between successive timesteps. As a result, we only had to convert the model so each successive timestep spanned one year. The default parameterizations in DICE-Climate and PAGE-Climate are specific to the non-annual timesteps. As a robustness check, we estimate new parameters that minimize the normalized root mean square error between output from the annual and default versions of each model. This procedure produces nearly identical temperature and sea level projections (Fig. 4.1), suggesting the annualization process did not drastically alter the default model behavior.

The model calibration period begins in 1850. However, FUND-Climate, DICE-Climate, and PAGE-Climate initialize in 1950, 2009, and 2005. We therefore change each model's initial conditions to 1850 values and run them using historic Representative Concentration Pathway (RCP) forcings and emissions. In PAGE-Climate, global temperature projections also depend on regional sulphate aerosol emissions (Hope 2011). For the period from 1850-2005, we aggregate national sulphate aerosol emissions (Smith et al. 2011) up to the PAGE regions and estimate each region's annual share of the global total (Fig. 4.2). We apportion global RCP sulphate emissions to each PAGE-Climate region based on these estimated shares. For periods beyond 2009, we assume global RCP emissions are portioned out according to the default regional emission shares from PAGE's business as usual scenario. Regional emission shares from 2005-2009 are linearly interpolated to ensure a smooth transition between the historical and projection periods.




Figure 4.1 | Annualized versions of DICE-Climate and PAGE-Climate produce similar output to the default versions. a, Global mean surface temperature increase projected by the default version of DICE-Climate that runs for ten-year timesteps (circles) and the annualized version (black line). The

parameters in the annualized version were calculated such that they minimized the normalized root mean square error between the output from both model versions. **b**, The same as in **a** but for global mean sea level. **c-d**, The same as in **a-b**, but for PAGE-Climate under its default irregularly spaced timesteps.




Figure 4.2 | Shares of historic global anthropogenic sulfur dioxide emissions for the eight PAGE-Climate regions. Each line shows the share of global anthropogenic sulphate aerosol emissions for the eight PAGE-Climate regions. Regional shares for 1850-2005 are calculated by aggregating country-level emission estimates from Smith et al. (2011). Emission shares after 2009 correspond to the default values from PAGE-Climate’s business-as-usual scenario, with 2005-2009 shares (light gray shaded region) linearly interpolated to enable a smooth transition between the two periods.

Calibration and Forcing Data

We calibrate SNEASY+BRICK and the three IAM climate models from 1850-2017 to observations of atmospheric carbon dioxide concentrations from Mauna Loa (Keeling et al. 1976; Tans and Keeling 2018) and Law Dome ice cores (Etheridge et al. 1996), global mean surface temperature (Morice et al. 2012), and global mean sea level changes (Church and White 2011). For SNEASY+BRICK, we also include chlorofluorocarbon based estimates of oceanic carbon fluxes, global ocean heat uptake between 0-3,000 meters (Gouretski and Koltermann 2007), and individual contributions to sea level from thermal expansion (Church et al. 2013), glaciers and small ice caps (Dyurgerov and Meier 2005), the Greenland ice sheet (Sasgen et al. 2012; The IMBIE Team 2020), and the Antarctic ice sheet (The IMBIE Team 2018). Our calibration framework accounts for time-varying observation measurement errors provided with each data set (described below).

Official U.S. SCC values are estimated using four high-emission scenarios and one low - emission scenario based on results from the 22nd Energy Modeling Forum exercise (Fig. 4.3).

The four high emission scenarios are designed such that atmospheric CO₂ concentrations range between 612-889 ppm in 2100 (Interagency Working Group on the Social Cost of Carbon 2010). For our baseline results, we run each model using greenhouse gas emissions and radiative forcing values from the high-emission, no-mitigation Representative Concentration Pathway 8.5 (RCP8.5) scenario (Meinshausen 2010; Meinshausen et al. 2011b). When forced by this scenario, SNEASY+BRICK projects a 2100 atmospheric CO₂ concentration of 907 ppm (873-947 ppm, 5-95% range). This is on the high end of, but generally consistent with, the U.S. scenario concentration range given the different model structures and carbon cycle uncertainties we consider. In a sensitivity test, we repeat our analysis using the RCP2.6 scenario. We also maintain each IAM's default socioeconomic scenarios.




Figure 4.3 | Comparison of RCP carbon dioxide emissions with the emission scenarios used for official U.S. social cost of carbon estimates. Solid lines show fossil and industrial CO₂ emissions for RCP2.6 (blue) and RCP8.5 (red). Colored circles show the four high-emission scenarios used in official U.S. SCC estimates for 2000, 2010, 2020, 2030, 2050 and 2100. Colored diamonds show values for the corresponding low-emissions U.S. scenario (Interagency Working Group on the Social Cost of Carbon 2016).

Model Calibration

We closely follow past work (Urban and Keller 2010; Ruckert et al. 2017; Wong et al. 2017; Vega-Westhoff et al. 2019; Errickson et al. 2020) and calibrate each model to the historic climate record using a Bayesian framework. The posterior probability p of the uncertain parameters θ gives the updated probability of observing them after assimilating the climate data y and is given by Bayes' theorem (Bayes and Price 1763). We sample the joint posterior parameter distribution using a robust adaptive Metropolis Markov chain Monte Carlo (MCMC) algorithm, which constructs a Markov chain of parameter samples that eventually converge to the target (posterior)

distribution (Metropolis et al. 1953; Hastings 1970; Gilks, Richardson, and Spiegelhalter 1996; Vihola 2012). The MCMC algorithm we use aids convergence to this target distribution by adapting how it explores the parameter space as the underlying posterior correlation structure begins to emerge. This calibration framework also requires calculating Bayes’ theorem up to proportionality:

$$p(\theta|y) \propto p(\theta)L(y|\theta), \quad (4.2)$$

where $p(\theta)$ represents the prior parameter probability and $L(y|\theta)$ is the likelihood function. We thus avoid using the numerical approximation techniques typically required for Bayesian inference (Smith and Roberts 1993; Gilks, Richardson, and Spiegelhalter 1996), rendering adaptive MCMC a computationally efficient calibration technique for the reduced-complexity models we consider.

We select physically informed prior parameter distributions for SNEASY+BRICK consistent with past model calibration exercises (Urban and Keller 2010; Wong et al. 2017; Vega-Westhoff et al. 2019; Errickson et al. 2020). Typically, BRICK’s Antarctic ice sheet (AIS) model first undergoes an initial, separate calibration to paleoclimate data spanning 240,000 years before the current era to the present. This procedure helps constrain deeply uncertain AIS dynamics, such as the potential for rapid ice sheet disintegration due to marine ice sheet instability and hydrofracturing, that are not represented in the modern observational record (Wong, Bakker, and Keller 2017). In this study, we instead form the AIS prior distributions for our calibration to modern observations using the parameter samples from a previous AIS paleoclimatic calibration. This involves fitting individual kernel density estimates (KDEs) for each uncertain AIS parameter to 800,000 paleoclimatic calibration samples (Table C.4). While using KDEs neglects potential AIS parameter correlations from the paleoclimate record, we obtain similar results with SNEASY+BRICK when using a multivariate normal distribution fit to the full AIS paleoclimatic calibration sample.

All models use the same equilibrium climate sensitivity (ECS) prior distribution that approximates the information in paleo observations that are not accounted for in our calibration (Urban and Keller 2010). We truncate this distribution from 0°C to 10°C, the same range spanned by the ECS distribution used for official U.S. SCC estimates (Interagency Working Group on the Social Cost of Carbon 2010). Except for several parameters describing initial conditions, we use truncated uniform prior distributions for the uncertain IAM climate model parameters. When available, we use an IAM’s default distribution to inform our prior ranges. PAGE-Climate uses triangular distributions as a default, so we set our prior ranges to be slightly larger than these triangular distribution bounds (Table C.3). In contrast, FUND-Climate uses untruncated gamma and normal default distributions for most of its uncertain parameters (Table C.2). We therefore set FUND-Climate’s prior bounds as the minimum and maximum values that occur in ten million samples from each of the default distributions. For parameters without default distributions, such as all those found in DICE-Climate (Table C.1), we generally select wide, physically-informed prior ranges that contain the model’s default parameter value.

We model the likelihood of the observations in terms of the data-model residuals:

$$r_{it} = y_t - f(\theta_{it}), \quad (4.3)$$

where $f(\theta_{it})$ represents the climate model output at time t for parameter sample i . This residual consists of the discrepancy between the observations and modeled climate ε as well as a time-varying observation error term ω :

$$r_{it} = \varepsilon_{it} + \omega_t. \quad (4.4)$$

We use a first order autoregressive-process (AR1) model that accounts for time-varying observation errors contained in the data (Ruckert et al. 2017) to approximate residuals for the global mean surface temperature anomaly, oceanic heat uptake, global mean sea level, as well as individual sea level contributions from glaciers and small ice caps, the Greenland ice sheet, and the Antarctic ice sheet. We adopt a similar approach for ice core and more recent instrumental atmospheric CO₂ concentration observations, but use a first-order continuous time autoregressive (CAR1) model to account for irregular spacing between CO₂ observations (Koen 2005; Errickson et al. 2020). We do not prespecify values for the AR1 and CAR1 statistical process parameters, but instead estimate them alongside the uncertain climate parameters during model calibration.

We initialize the MCMC algorithm using the mean posterior parameter values and covariance matrix from previous model calibration exercises to begin sampling in a region of high posterior probability. We create a Markov chain of six million parameter samples and discard the first two million iterations to reduce the impacts of a potential burn-in. The remaining chain is thinned at equally spaced intervals to produce 10,000 posterior parameter samples that we then use to for probabilistic climate projections and SCC estimates.

Model Testing and Diagnostics

We carry out probabilistic global surface temperature and sea level prediction tests in addition to a range of model diagnostics for DICE-Climate, FUND-Climate, PAGE-Climate, and SNEASY+BRICK. We use three approaches to explore how different treatments of parametric uncertainties affect climate model behavior: a baseline case using posterior parameter samples from our Bayesian calibration framework (“Baseline”), a case that removes the posterior correlations by independently sampling each marginal posterior parameter distribution (“Remove Correlations”), and a case reflecting the uncertainty framework used for official U.S. SCC estimates that neglects parameter correlations and samples a prespecified equilibrium climate sensitivity (ECS) distribution (“U.S. Uncertainty”).

Out-of-sample probabilistic prediction tests

We calibrate each climate model following the Bayesian framework described above using a reduced set of observations covering the 1850-1970 period. For each of the 10,000 posterior parameter samples, we project global mean surface temperature and sea level from 1850 to 2017. The models are evaluated based on how skillfully they predict the observations and their associated uncertainty levels during the 1971-2017 out-of-sample period. We also explore how

sensitive each model is to the amount of data used in the calibration by comparing the out-of-sample predictions to those made using parameters calibrated to the full set of observations spanning 1850-2017. Predictions are also made for the “Remove Correlations” scenario to explore how parameter interactions affect out-of-sample uncertainty levels.

100 GtC impulse response behavior

We carry out a suite of impulse response tests relevant for assessing climate model performance (Joos et al. 2013; Schwarber et al. 2019) and estimating the SCC (National Academies of Sciences, Engineering, and Medicine 2017). For each posterior parameter sample, we perform two climate model runs. The first run represents a reference case while the second run adds a 100 GtC emissions pulse in 2020 but is otherwise identical. Although past impulse response tests are often carried out against a constant background CO₂ concentration level (Joos et al. 2013), we conduct this analysis for the RCP2.6 and RCP8.5 scenarios. This enables us to compare how models respond to a more realistic and policy-relevant setting where background climate conditions change over time (National Academies of Sciences, Engineering, and Medicine 2017; Schwarber et al. 2019). We carry out these impulse tests for all three parametric uncertainty frameworks listed above. For the “U.S. Uncertainty” impulse response tests, we sample each IAM climate model’s default parameter distributions as well as the ECS distribution used for official U.S. SCC estimates (Interagency Working Group on the Social Cost of Carbon 2010). Under this framework, we run DICE-Climate and PAGE-Climate using 10-year time steps to capture the model behavior that influenced past SCC estimates (we make this simplifying assumption for PAGE-Climate because we do not have an irregular timestep spacing scenario that covers the full 1850-2300 projection period, but note the default version of the model uses ten-year timesteps between 2010-2050 when the pulse occurs). We calculate five impulse response metrics as the difference between the 100 GtC pulse and reference model run:

- (i) The evolution of the annual change in global mean surface temperature over time.
- (ii) The evolution of the annual change in global mean sea level over time.
- (iii) The peak temperature that occurs between 2020-2300.
- (iv) The number of years required to achieve the peak temperature increase.
- (v) The increase in global mean sea level occurring one hundred years after the emission pulse.

Temperature response to cumulative CO₂ emissions

We test the relationship between the increase in global mean surface temperature and cumulative CO₂ emissions. Past work using comprehensive Earth system models shows a near linear increase in temperature with cumulative emissions up to at least 2,000 GtC (Matthews et al. 2009; National Academies of Sciences, Engineering, and Medicine 2017). We assess this relationship using projections where CO₂ concentrations increase 1% per year and all other external forcings remain fixed at their 1850 values. For each posterior parameter sample, we first calculate a time series of CO₂ concentrations that increase by 1% per year starting from that sample’s uncertain initial atmospheric CO₂ concentration. We then calculate the CO₂ emissions time series that minimizes the root-mean-square error (RMSE) between modeled atmospheric CO₂ concentrations and the 1% per year increasing concentration scenario. The climate model is

then forced with this inverted emissions series to project future temperatures. We repeat this process using 1,000 randomly selected posterior parameter samples.

Estimating the Social Cost of Carbon

We couple SNEASY+BRICK to the non-climate components of the DICE, FUND, and PAGE IAMs. We use these updated models to estimate the SCC and refer to them as *SB-DICE*, *SB-FUND*, and *SB-PAGE* throughout the manuscript.

The default version of PAGE calculates climate damages based on local temperatures that vary based on regional sulphate aerosol emission levels. Because SNEASY+BRICK does not project regional temperatures, we employ a pattern scaling technique that allows us to couple SNEASY+BRICK with the PAGE IAM (Mitchell 2003; Lynch et al. 2017). We estimate regional temperatures $T_{Regional}$ based on the global average surface temperature projected by SNEASY+BRICK T_{Global} using:

$$T_{Regional,t} = \beta_0 + \beta_1 T_{Global,t} + \beta_2 T_{Global,t}^2 + \varepsilon_t . \quad (4.5)$$

For each PAGE region, we estimate Equation (S.5)'s coefficients by using a linear regression on the regional and global temperatures from 100,000 runs of the default PAGE climate model. The resulting pattern scaled temperatures (Fig. 4.4a) lead to near identical SCC distributions in the default version of PAGE (Fig. 4.4b).




Figure 4.4 | Temperature pattern scaling in PAGE. **a**, Comparison of European Union temperature projections for the default PAGE climate model and the version using pattern scaling. Results correspond to projections using 5000 samples from PAGE's default climate parameter distributions. The diagonal red line shows where the points would fall if the default and pattern scaled projections were identical. **b**, SCC distribution using the default version of PAGE (blue) and the version that projections regional temperatures using pattern scaling (red). The vertical lines show the 95th percentile SCC estimate for each

model version. SCC estimates are for a CO₂ emission pulse in 2009 using a constant 3% consumption discount rate.

Baseline estimate

To estimate the SCC, we produce two climate projections to 2300 for the RCP8.5 scenario; a “baseline” projection and a “pulse” projection that adds a one GtC emission pulse in 2020. These projections contain a time-series of simulated noise calculated from the calibrated statistical process parameters. This noise accounts for structural model errors, observation errors, and short-term natural variability that is not resolved by SNEASY+BRICK (Ruckert et al. 2017). We force each IAM with these climate projections and produce two climate change damage trajectories. The SCC for posterior parameter sample i is then calculated as the discounted sum of the difference in damages between the “baseline” and “pulse” model runs:

$$SCC_i = \sum_{t=0}^T \sum_{r=1}^R (Damages_{itr}^{pulse} - Damages_{itr}^{baseline}) \times \frac{1}{(1 + \rho)^t} . \quad (4.6)$$

Here, t indexes time, r indexes the IAM region, and ρ is a constant consumption discount rate.

We normalize our SCC estimates to the emission pulse size and express them in 2007 U.S. dollars per ton of emitted CO₂ to facilitate comparison with official U.S. SCC estimates (Interagency Working Group on the Social Cost of Carbon 2010; United States Department of Labor: Bureau of Labor Statistics 2020). Our main estimates use a constant 3% discount rate, but we also provide additional estimates using different discount rates as well as a Ramsey discounting framework. For SB-DICE, we linearly interpolate SCC estimates for 2015 and 2025 to produce our 2020 estimates (Interagency Working Group on the Social Cost of Carbon 2010). For SB-PAGE, we assume values between 2200-2300 remain fixed at their 2200 values.

When we report uncertainty ranges for our SCC distributions, we provide the 5th and 95th percentile credible intervals. For multi-model results, we assume each IAM has an equal probability and average across the individual model estimates.

Alternative estimates

We provide additional SCC estimates to explore how different uncertainty frameworks and modeling assumptions affect distributions of the SCC. These analyses include:

- (i) Neglecting the posterior climate parameter correlations that emerge during model calibration.
- (ii) Neglecting the posterior climate parameter correlations that emerge during model calibration and sampling the prespecified ECS distribution used for official U.S. SCC estimates. We specify this distribution as $\frac{1.2}{1-f}$ where $f \sim \mathcal{N}(0.62, 0.034)$ and

the ECS is truncated between 0 and 10 °C (Marten 2011; Interagency Working Group on the Social Cost of Carbon 2010).

- (iii) Using climate projections based on parameters from a calibration of SNEASY+BRICK that does not include any sea level observational constraints.
- (iv) The RCP2.6 emissions and forcing scenario.
- (v) Using constant consumption discount rates (2.5%, 5%) and Ramsey discounting (pure rate of time preference = 1.5%, elasticity of marginal utility of consumption = 1.0 and 1.5)
- (vi) An equity-weighting framework.

Equity-weighting

Equity-weighted SCC estimates assume an additional dollar of consumption improves the welfare of a poor person more than a rich person and have been thoroughly described elsewhere (Anthoff, Hepburn, and Tol 2009; Anthoff and Emmerling 2019). We use a standard utilitarian social welfare function and estimate the equity-weighted SCC for posterior parameter sample i with:

$$SCC_i^{equity} = \sum_{t=0}^T \sum_{r=1}^R \left(\frac{c_{i0r}}{c_{i0x}} \right)^{-\eta} \left(\frac{c_{itr}}{c_{i0r}} \right)^{-\eta} (1 + \rho)^{-t} (Damages_{itr}^{pulse} - Damages_{itr}^{baseline}). \quad (4.7)$$

Here η is the consumption elasticity of marginal utility, c represents regional per capita consumption levels, ρ is the pure rate of time preference, and x corresponds to the region used to normalize the equity-weighted SCC into a welfare-equivalent consumption loss. We set ρ to 1.5%, use η values of 1.0 and 1.5, and report results using the United States, Africa, and Latin America as normalization regions. Because DICE is a global model and cannot produce equity-weighted SCC estimates, we use its regional counterpart the Regional Integrated Climate Economy (RICE) model (Nordhaus 2010).

4.3 Variation in Climate Hindcasts and Impulse Response Behaviors

The ability to approximate the historic climate record reasonably well is a crucial test to assess model-based projections (Flato et al. 2013). The DICE, FUND, and PAGE climate models with the default parameter distributions used in official U.S. SCC estimates arguably fail this test as they produce rather poor temperature and sea level hindcasts (Fig. C.1). DICE-Climate is rather sensitive to strengthening greenhouse gas forcing, causing it to overpredict observed warming from the 1950s onwards (Fig. C.1a). FUND-Climate's hindcast is generally consistent with the temperature trend, but also overconfident (Fig. C.1b). It produces a 5-95% confidence interval narrower than the reported range of temperature observation errors. PAGE-Climate also slightly overestimates the temperature trend and produces rapidly increasing uncertainty levels towards

the end of the hindcast period (Fig. C.1c). For global mean sea level, each model again struggles to replicate the historic trend and observed levels of uncertainty (Fig. C.1d-f).

Recalibrating each climate model (see Methods) substantially changes some of the original parameter distributions, suggesting the original parameter values may be inconsistent with the historic climate record (Fig. C.2). For instance, FUND-Climate's posterior distribution for the parameter governing the temperature sensitivity of terrestrial carbon stocks shifts towards lower values and experiences a 78% reduction in its 95th percentile estimate (Fig. C.2b). Our calibrated IAM ECS distributions are on the low end of likely values (Knutti and Hegerl 2008; Collins et al. 2013; Sherwood et al. 2020) and differ substantially from the ECS distribution used for official SCC estimates (Fig. C.2a). The U.S. ECS distribution is parameterized to likely ECS ranges reported in the Intergovernmental Panel on Climate Change's (IPCC) Fourth Assessment Report, resulting in a mean of 3.5 °C and a 95th percentile value of 7.1 °C (Interagency Working Group on the Social Cost of Carbon 2010). In contrast, the mean IAM posterior ECS values range between 2.2-2.3 °C, with 95% credible interval estimates spanning 3.6-3.9 °C.

SNEASY+BRICK's posterior ECS estimate of 3.1 °C (2.1-4.3 °C, 5-95% range) is slightly narrower but generally consistent with recent ECS estimates that jointly account for information on climate system feedbacks, the paleoclimate record, and recent warming levels (Fig. C.2a) (Sherwood et al. 2020).

Correlations also emerge between the uncertain climate model parameters that past impulse response evaluations and SCC estimates do not account for (Fig. C.3). In FUND-Climate, the sea level response to temperature shows a strong positive correlation with the sea level e-folding timescale ($corr. = 0.94$), suggesting stronger temperature sensitivities can be compensated for by slowing the rate at which sea level equilibrates (Fig. C.3d). Similarly, in PAGE-Climate a negative correlation ($corr. = -0.82$), emerges between the parameters affecting direct and indirect sulphate aerosol forcing levels (Fig. C.3f). These and other correlations lend evidence to the idea that past SCC estimates suffer from two sources of parameter-related bias: (i) the use of parameter distributions inconsistent with past observations and (ii) the effects of parameter combinations that are inconsistent with observations caused by neglecting their underlying correlation structure. Below, we show how these correlations influence climate model impulse response behaviors and can act as an important constraint on the SCC's upper tail.

We first evaluate each recalibrated climate model through probabilistic out-of-sample predictions of global surface temperature and sea level changes from 1970-2017 (Fig. 4.5, Fig. C.4). SNEASY+BRICK, and to a slightly lesser degree, FUND-Climate project the mean temperature trend and its observed uncertainty range reasonably well (Fig. 4.5b and d). Although DICE-Climate and PAGE-Climate produce a skillful mean response when each model is calibrated to the full set of observations, their out-of-sample predictions deviate from the historic temperature trend (Fig. 4.5a and c). Across all models, neglecting the posterior parameter correlations produces much larger uncertainty levels. For instance, removing the correlations increases the 95% predictive temperature credible interval in 2017 by 17-44%. Similar patterns again emerge for the global mean sea level hindcasts, with the posterior correlations acting as an even stronger constraint to prevent large, unrealistic uncertainty ranges during the out-of-sample period (Fig. C.4).




Figure 4.5 | Climate model structure and parameter correlations influence out-of-sample temperature prediction skill. Predicted global mean surface temperature anomalies relative to the 1861-1880 mean using the DICE-Climate (a, light orange), FUND-Climate (b, green), PAGE-Climate (c, blue), and SNEASY+BRICK (d, gold) models. Each model is calibrated to observations covering the period 1850-1970 and then performs an out-of-sample prediction test from 1970-2017 (gray shaded region). Center lines outlined in black depict the mean model response for 10,000 posterior parameter samples and the edges of the colored regions span the 95% predictive credible intervals. The dashed black center line shows the mean model response when calibrating each model to the full set of observations from 1850-2017. Purple lines show the 95% predictive credible intervals for a scenario that neglects posterior parameter correlations by independently sampling each posterior parameter distribution. The solid red line depicts temperature observations with lighter shading for the 2σ observation error range.

The temperature response of each IAM's climate model following a 100 GtC emissions pulse also exhibits considerable sensitivity to a refined model calibration. (Fig. 4.6). After calibration, the IAM climate model temperature responses differ markedly from those made with the default

parameter distributions used in official SCC estimates (Tables C.1-3) (Interagency Working Group on the Social Cost of Carbon 2016). Their impulse response behaviors also show visibly decreased uncertainty levels, however this uncertainty increases substantially after removing the parameter correlations. As discussed above, the IAM climate models have been criticized for having too slow of a temperature response timescale (National Academies of Sciences, Engineering, and Medicine 2017; Dietz et al. 2020). After calibration DICE-Climate and FUND-Climate both show a more rapid response time and an even quicker drawdown from the peak temperature. While these results highlight how previously neglected parametric uncertainties can strongly influence impulse response behaviors, they also show that thoroughly calibrating the IAM climate models cannot fully overcome the deficiencies in their underlying structure. We obtain similar results for other impulse response metrics important for estimating the SCC. Removing parameter correlations or adopting the U.S. uncertainty framework leads to increased uncertainty in the peak temperature arising from the 100 GtC emissions pulse (Fig. C.5a-d), the years it takes to achieve this peak temperature (Fig. C.5e-h), and the relationship between total warming and cumulative CO₂ emissions (Fig. C.6).




Figure 4.6 | Temperature response dynamics and uncertainty levels resulting from a 100 GtC emission pulse vary with climate model parameter choice. Predicted global mean surface temperature increase from a 100 GtC emission pulse in 2020 under the RCP8.5 scenario using the DICE (a, orange),

FUND (b, green), PAGE (c, blue), and SNEASY+BRICK (d, gold) climate models. The lines corresponding to each model's specific color depict 1,000 randomly sampled runs from a 10,000 member ensemble using the posterior parameter samples, with the center white line outlined in black showing the mean impulse response for the full ensemble. Purple lines similarly show 1,000 randomly sampled runs from a 10,000 member ensemble that neglects the posterior parameter correlations by sampling the marginal posterior distributions independently. For the DICE, FUND, and PAGE climate models, the solid red line shows the mean model response from 10,000 runs sampling the models' default, uncorrelated parameter distributions and the equilibrium climate sensitivity distribution used for official U.S. SCC estimates. Under this scenario, the DICE and PAGE models run on ten-year time steps. Dashed red lines show the 95% confidence interval. For SNEASY+BRICK, the red lines depict a similar scenario that neglects posterior correlations and also samples the U.S. equilibrium climate sensitivity distribution.

Although SCC estimates incorporate sea level-related damages, the sea level impulse response behaviors have received rather scarce attention in the literature. Each IAM climate model projects a global mean sea level response that is substantially lower than SNEASY+BRICK (Fig. C.7). A century after the 100 GtC emissions pulse, DICE-Climate, FUND-Climate, and PAGE-Climate project approximately 0.03-0.04 m of sea level rise (Fig. C.5i-k) while SNEASY+BRCK projects a nearly 0.12 m (0.08-0.16, 5-95% range) increase (Fig. C.5l). In fact, the IAM climate models' 95% percentile values are still lower than SNEASY+BRICK's mean response, ranging between 0.04-0.11 m. SNEASY+BRICK's higher sea level impulse values are driven, in part, by the model's improved, physically-motivated representation of individual sea level contributors, including its representation of deeply uncertainty ice sheet dynamics (Wong et al. 2017; Wong, Bakker, and Keller 2017). Under certain parameter settings, for example, the 100 GtC emissions pulse produces accelerated sea level rise by triggering a rapid disintegration of the Antarctic ice sheet that the other models do not account for.

4.4 Correlations Constrain Upper Tail SCC Estimates

Coupling the SNEASY+BRICK climate model to each IAM, we estimate a multi-model mean SCC of \$49/t-CO₂ (\$26-65/t-CO₂, 5-95% range) for an emission pulse in 2020 under the RCP8.5 scenario and a constant 3% consumption discount rate (Fig. 4.7, Table C.5). We obtain similar results using a Ramsey discounting framework that yields declining effective discount rates with slowing economic growth (Fig. C.8, Table C.5) (Arrow et al. 2014; National Academies of Sciences, Engineering, and Medicine 2017). Although official U.S. SCC estimates are based on different socioeconomic and emission scenarios than the ones we consider (see Methods), a comparison with our results still proves informative. Relative to our central estimates, the average U.S. SCC is 14% lower while the corresponding 95th percentile estimate is 89% higher (Interagency Working Group on the Social Cost of Carbon 2016). This is broadly consistent with the expected effects of more realistic temperature dynamics increasing the SCC (Dietz et al. 2020; Hänsel et al. 2020) and our calibrated ECS distribution helping constrain extreme SCC values (Errickson et al. 2020).




Figure 4.7 | Social cost of carbon uncertainty and 95th percentile estimates vary with treatment of parametric uncertainties. SCC distributions estimated with the SB-DICE (a, orange), SB-FUND (b, green) and SB-PAGE (c, blue) integrated assessment models under the RCP8.5 scenario, with vertical colored lines highlighting 95th percentile SCC estimates. Distributions with each model’s specific color depict baseline SCC estimates using SNEASY+BRICK’s calibrated parameter samples (“Baseline Estimate”). Purple distributions show SCC estimates based on climate projections that neglect the posterior correlations found in the “Baseline Estimate” scenario (“Remove Correlations”). Red lines depict SCC estimates for a scenario that neglects posterior parameter correlations and samples the equilibrium climate sensitivity distribution specified for official U.S. SCC estimates (“U.S. Uncertainty”). The dashed black line depicts the estimated SCC distribution for a scenario that recalibrates SNEASY+BRICK without sea level observations (“No Sea Level Data”). Each distribution corresponds to 10,000 SCC estimates for an emission pulse in 2020 using a constant 3% consumption discount rate.

While neglecting the climate model parameter correlations (see Methods) produces almost no change in the multi-model mean SCC – causing it to increase by less than 1% – the 95th percentile estimate increases by 23% (Fig. 4.7, Table C.5). As discussed above, ignoring these correlations substantially increases uncertainty in temperature and sea level impulse response behaviors, producing more extreme outcomes for each (Fig. 4.6, Fig. C.7). Parameter correlations also strongly influence climate projection uncertainties. Without them, the 95th percentile temperature projection in 2100 increases by 0.5 °C, while the corresponding 95th percentile sea level increases by more than a meter (Fig. C.10). As a result of their influence on extremes in the climate system, these parameter correlations also impact SCC extremes. For instance, the positive correlation between aerosol cooling and the ECS constrains the SCC’s upper tail by offsetting some of the warming that would otherwise occur under high ECS values (Fig. 4.8).

The SCC exhibits a marked shift towards higher values under the U.S. uncertainty framework (Fig. 4.7, Table C.5). Official U.S. SCC estimates sample each IAM’s default uncertain parameters (Tables C.1-3), but also prescribe the same ECS distribution across all models (Interagency Working Group on the Social Cost of Carbon 2016). As described above, this distribution is inconsistent with SNEASY+BRICK’s posterior ECS distribution that is derived from the observations as well as recent work showing a reduced probability for ECS values above approximately 6 °C (Fig. C.2a) (Sherwood et al. 2020). Removing the correlations in SNEASY+BRICK’s calibrated parameters and pairing them with the U.S. ECS distribution (see Methods) increases the SCC by 8%. The corresponding 95th percentile estimate increases by 92% and now better aligns with the official U.S. SCC values. To illustrate the magnitude of this change, if we apply the lowest constant discount rate used for U.S. SCC estimates (2.5%) to our

baseline results with fully correlated parameters, the 95th percentile SCC increases by 62% (Table C.5). Using the highest discount rate (5%) yields a 74% decrease. As these results illustrate, different treatments of the uncertain climate parameters can produce changes in the SCC's upper tail on par with the effects of discount rate choice.




Figure 4.8 | Climate model calibration and parameter correlations constrain high social cost of carbon estimates. Relationship between the equilibrium climate sensitivity and the SCC estimated with the SB-DICE (a, orange), SB-FUND (b, green) and SB-PAGE (c, blue) integrated assessment models under the RCP2.6 (triangles) and RCP8.5 (circles) scenarios. For the baseline RCP8.5 and RCP2.6 SCC estimates, each model's specific color corresponds to an uncertain parameter that scales aerosol radiative forcing (with white signaling weaker aerosol cooling). Purple circles depict SCC estimates for RCP8.5 using climate model parameter samples that neglect parameter correlations by sampling each posterior parameter distribution independently. Gray circles similarly show SCC estimates that neglect parameter correlations while also sampling the equilibrium climate sensitivity distribution used for official U.S. SCC estimates. All SCC estimates correspond to an emission pulse in 2020 using a constant 3% consumption discount rate, with 2,500 randomly sampled estimates plotted for each scenario.

4.5 Sensitivity Analyses and Discussion

Incorporating sea level observations into the climate model calibration affects the SCC in two main ways: via changes in sea level and changes in temperature projections. The unconstrained sea level projections are higher than their baseline counterpart (Fig. C.10b), causing an increase in estimated sea level-related damages. Consistent with past work (Vega-Westhoff et al. 2019), accounting for sea level observational constraints in the model calibration also influences the posterior ECS distribution (Fig. C.2a) and resulting temperature projections (Figure C.10a). The ECS produced without sea level data is 2.8 °C (1.9–4.1 °C, 5–95% range), an approximate 10% reduction in the mean relative to our baseline estimate of 3.1 °C (2.1–4.3 °C, 5–95% range). The indirect temperature effect caused by these lower ECS values and the direct sea level damages effect influence the SCC in opposite ways. Yet partly because temperature impacts are the primary driver of each IAM's damage functions (Diaz and Moore 2017; National Academies of

Sciences, Engineering, and Medicine 2017), the indirect effect dominates and the mean SCC declines by approximately 6% (Fig. 4.7, Table C.5).

Several climate model impulse response behaviors as well as our SCC estimates are sensitive to assumptions about future emission scenarios (Fig. C.5, Fig. C.10). Under RCP2.6's less severe climate forcing, SNEASY+BRICK projects 0.08 m (0.04-0.1 m, 5-95% range) of additional sea level rise a century after a 100 GtC emissions pulse (Fig. C.5l). This represents a 33% reduction relative to SNEASY+BRICK's mean RCP8.5 response and stands in stark contrast to the IAM climate models, which exhibit less sensitivity to the choice of forcing scenario (Fig. C.5i-k). Consistent with past analyses (National Academies of Sciences, Engineering, and Medicine 2017), each model's temperature impulse also draws down less rapidly from its peak value because it no longer occurs against the backdrop of RCP8.5's accelerated warming levels (Fig. C.9). The resulting RCP2.6 SCC of \$16/t-CO₂ (\$9-24/t-CO₂) is 67% lower than the corresponding RCP8.5 estimate. Perhaps surprisingly, the SCC's upper tails are more sensitive to parameter correlations under the RCP2.6 scenario (Fig. C.10c-e). Their removal increases the 95th percentile SCC by more than 70%.

Equity-weighting provides an additional framework for estimating the SCC that assigns greater value to climate impacts experienced by the poor (Anthoff, Hepburn, and Tol 2009). This framework accounts for distributional equity concerns and normative ethical preferences by formalizing the concept that an extra dollar of consumption improves the well-being of a poor person more than a rich one. As a result, equity-weighted SCC estimates vary based on the socioeconomic status of the region they are being estimated for and may even differ by more than an order of magnitude between low- and high-consumption regions. Our central equity-weighted SCC estimates (see Methods) are \$11/t-CO₂ (\$6-15/t-CO₂, 5-95% range) for Africa, \$43/t-CO₂ (\$23-58/t-CO₂, 5-95% range) for Latin America, and \$185/t-CO₂ (\$97-260/t-CO₂, 5-95% range) for the United States. Across the models, regions, and levels of inequality aversion considered, we uncover a robust response in how equity-weighted SCC upper tails vary with the treatment of parametric uncertainties (Fig. C.11). Without correlations, our 95th percentile equity-weighted SCC for Africa, Latin America, and the United States increases to \$18/t-CO₂, \$70/t-CO₂, and \$312/t-CO₂. Under the U.S. uncertainty framework, these values increase further still to \$29/t-CO₂, \$113/t-CO₂, and \$482/t-CO₂.

4.6 Conclusion

Our SCC estimates are subject to several caveats that suggest avenues for future research. While the SNEASY+BRICK model improves on the representation of the climate system in DICE, PAGE, and FUND, it does not incorporate the full range of feedbacks and complex spatial patterns found in more sophisticated Earth system models (Flato et al. 2013; Voldoire et al. 2019). Although we approximate the effects of one climate tipping point, a rapid Antarctic ice sheet disintegrations, other deeply uncertain threshold responses in the climate system (Lenton et al. 2019) would likely influence our estimates of the SCC's upper tail. We also use the default socioeconomic components of DICE, FUND and PAGE, thereby neglecting recent empirical work on climate damage functions (Carleton and Hsiang 2016) and the potential for compound extreme climate events (Zscheischler et al. 2018). Recent work also calls into question whether current IAM damage functions appropriately capture the risk society faces from sea level-related

impacts such as extreme flooding (Kirezci et al. 2020), enhanced erosion of the world's coastlines (Vousdoukas et al. 2020), and the heightened vulnerability of disadvantaged populations to coastal storms (Green et al. 2010; Dasgupta et al. 2011; Lieberman-Cribbin et al. 2020). Relatedly, exploring the joint effect of climate and socioeconomic uncertainties (Lamontagne et al. 2019), including links between poverty and increased climate risks that go beyond equity-weighting considerations (Dennig et al. 2015; Rao et al. 2017), would further enrich our analysis.

Upper tail SCC estimates are meant to capture low-probability, high-impact climate change outcomes in a simple policy metric. These estimates serve as an important tool for assessing the potential benefits of different policy proposals and more broadly inform the debate over what climate risk mitigation strategies society should pursue. We show that accounting for the joint intersection between carbon cycle, climate, and sea level parametric uncertainties acts as an important constraint on the SCC's deeply uncertain upper tail. Neglecting parameter correlations and sampling values inconsistent with the historic climate record substantially increases the 95th percentile SCC; a change comparable in magnitude to using the different discount rates considered for official U.S. SCC estimates.

CHAPTER 5

Conclusion

The design of effective and ethically defensible climate policies involves a number of challenges, to say the least. Deep uncertainties concerning climate system feedbacks and the potential for high-impact, irreversible threshold responses are unlikely to be fully resolved in the near future. At the same time, global inequality levels require difficult ethical judgments over how to value climate impacts experienced by different income groups across generations. In this dissertation, I used integrated assessment modeling frameworks to investigate how these challenges influence optimal climate policy and key metrics such as the social cost of carbon.

In Chapter 2, I quantify how the health co-benefits of CO₂ mitigation and co-harms of lost aerosol cooling affect optimal policies that target shared emissions sources. This analysis required the development of a new model, RICE+AIR, that accounted for the trade-off between CO₂ mitigation costs, avoided climate damages, and the climate and localized health consequences from air pollutant co-emissions. I valued the additional life years gained from improvements in local air quality using a utilitarian social welfare function that accounts for societal preferences on inequality. Across a range of robustness checks and different valuations of human health, accounting for co-benefits leads to increased CO₂ mitigation levels even with the additional harms from lost-aerosol cooling. In fact, assigning a VSL-like monetization to improvements in local health outcomes produced optimal peak temperatures consistent with the Paris Agreement 2.0 °C target. Out of concern that this modeling framework may be forcing regions to inefficiently use climate policy as a means to reduce air pollution, I also created a co-optimizing version of RICE+AIR that allows regions to mitigate CO₂ while simultaneously reducing SO₂, PM_{2.5}, and NO_x through independent end-of-pipe measures. Under this co-optimizing framework, accounting for health co-benefits still leads to increased CO₂ mitigation relative to the reference case that neglects them.

In Chapter 3, I provided improved SC-CH₄ estimates using an ensemble of IAMs and coupled climate, carbon, and methane cycle models. Importantly, these models accounted for the recent 25% upward revision to methane radiative forcing calculations that past SC-CH₄ studies overlooked. I calibrated each model in a Bayesian framework that samples the tails of the joint posterior parameter distribution and used these parameter samples to provide 100,000 SC-CH₄ estimates per model. The multi-model mean estimate for the SC-CH₄ was \$922/t-CH₄ (\$468-1540/t-CH₄, 5-95% range) using a constant 3% consumption discount rate under RCP8.5. Despite accounting for the increase in methane radiative forcing, this estimate is 23% lower than the central SC-CH₄ value used by the U.S. federal government. The corresponding 95th percentile SC-CH₄ estimate, used to account for high-impact climate scenarios in federal regulatory impact analyses, was more than 50% lower. These reductions occurred for two reasons: (i) the upper tail of the calibrated ECS distribution was better constrained relative to the distribution used by the U.S. government and (ii) posterior parameter correlations that emerge from the calibration prevent physically-implausible parameter combinations and extreme SC-

CH₄ estimates. I also provided the first probabilistic equity-weighted SC-CH₄ estimates. Equity-weighting assigns equal importance to the well-being of different individuals, but also represents the concept that an additional dollar increases the wellbeing of a poor person more than a rich person through a normative parameter capturing society's aversion to inequality. The central equity-weighted estimate for the United States was \$8040/t-CH₄ (\$4430-12,300/t-CH₄, 95% range) while sub-Saharan Africa's was \$130/t-CH₄ (\$72-199/t-CH₄, 95% range). However, higher values of the inequality aversion parameter described above increased the United States' SC-CH₄ to \$35,300/t-CH₄ (\$21,000-52,300/t-CH₄, 95% range) and decreased sub-Saharan Africa's to \$73/t-CH₄ (\$43-107/t-CH₄, 95% range). As these results illustrate, concerns for equity can substantially influence how society values climate impacts.

In Chapter 4, I quantified how different climate model calibration approaches and treatments of parametric uncertainties affect the SCC's decision-relevant upper tails. I first tested the ability of the simple climate models from the DICE, FUND, and PAGE IAMs to hindcast the historic climate record and produce impulse response behaviors consistent with more sophisticated Earth system models. This testing process involved two versions of each model, a baseline version sampling that model's default parameter distributions and an updated version that had been calibrated in a Bayesian framework. I found that both model versions performed rather poorly, suggesting a potential climate-related bias in past SCC estimates. While calibrating the IAM climate models did have a strong influence on their impulse response behaviors and uncertainty levels, it could not fully overcome each model's structural deficiencies. I then coupled a reduced-complexity Earth system model, SNEASY+BRICK, to the non-climate components of DICE, FUND, and PAGE to provide updated SCC estimates. My baseline mean SCC estimate was \$49/t-CO₂ (\$26-65/t-CO₂, 5-95% range) for a CO₂ emission pulse in 2020 under the RCP8.5 scenario. In a sensitivity analysis, I removed the posterior parameter correlations that emerged during model calibration and found the 95th percentile SCC estimate increased by 23% to \$80/t-CO₂. Under the U.S. climate uncertainty framework that neglects parameter correlations and samples a prespecified ECS distribution, the 95th percentile estimate increased further still to \$80/t-CO₂. I then replicated this analysis in an equity-weighting framework, uncovering a robust response in how the equity-weighted SCC upper tails varied with the treatment of parametric uncertainties. In the baseline case, the equity-weighted 95th percentile SCC estimates were \$15/t-CO₂ for Africa, \$58/t-CO₂ for Latin America, and \$260/t-CO₂ for the United States. Under the uncertainty framework used for official U.S. SCC estimates, these values nearly doubled to \$29/t-CO₂, \$113/t-CO₂, and \$482/t-CO₂.

While the three chapters of my dissertation provide new insights into the link between uncertainty, inequality, and global climate policy, they share several important caveats. First, each relies on simple, global climate models that do not capture climate change's complex spatial patterns. They also use the default IAM climate damage functions, which overlook recent, spatially-resolved empirical work showing large temperature effects on economic growth (Burke, Hsiang, and Miguel 2015), mortality (Carleton et al. 2020), agricultural productivity (Moore et al. 2017), labor supply (Zivin and Neidell 2014), and a range of other economic sectors (Carleton and Hsiang 2016). The combined effect of using simple, global climate models

and outdated damage functions produces results that neglect the fine-scale spatial heterogeneity of climate change impacts (Hsiang et al. 2017). As a result, I likely aggregate over important sources of climate-related inequality occurring at sub-regional scales (Scovronick, Budolfson, et al. 2019). Relatedly, while each chapter focuses on some aspect of equity and climate policy, I do not consider nonutilitarian ethical frameworks (Adler et al. 2017; Fleurbaey et al. 2019) and thus neglect an important source of ethical uncertainty.

Broadly, this dissertation shows that climate uncertainty and equity concerns can meaningfully impact the decision about how society should balance the upfront costs of greenhouse gas mitigation with the benefits of avoided damages. The caveats discussed above represent promising research avenues to expand upon these findings. In particular, future work should explore how the joint interactions between climate, socioeconomic, and ethical uncertainties not considered here affect policies designed to help those most vulnerable to a changing climate.

APPENDIX A

Supplementary Information for Chapter 2

A.1 Supplementary Figures

Note: “Reference case” refers to the model optimum where health benefits are excluded; aerosol co-reductions from CO₂ mitigation do occur, but only their climate impacts are included in the optimization. The reference case is therefore representative of standard cost-benefit climate-economy models. In contrast, “RICE+AIR” refers to the full model optimum that includes all climate and health impacts.




Figure A.1: Model comparison: RICE2010, Reference case, and RICE+AIR. The figure displays four optimal temperature trajectories that result from successive modifications to the unmodified, open-access Excel version of the RICE2010 model, culminating in the full RICE+AIR model that provides our main result:

1. RICE2010, which is the unmodified open-access Excel version of the model.
2. RICE2010 modified, which updates the population projections and modifies the social welfare function of RICE2010 (described in the Methods section) but maintains the exogenous aerosol assumptions.
3. Reference case, which adds endogenous aerosols but not health co-benefits to RICE2010 modified.
4. RICE+AIR, which adds health co-benefits to the Reference case and is the full model behind our main results.




Figure A.2: Optimal carbon prices. Carbon prices associated with the optimal decarbonization and temperature trajectories reported in Figure 2.2 of the main text. (Note: full decarbonization occurs at the point where the carbon prices peak and subsequently decline along a single line, representing a backstop technology, which is a technology that can replace all fossil fuels).




Figure A.3: Optimal decarbonization rates with a 3% discount rate. $\rho = 0, \eta = 1.3$.




Figure A.4: Emission intensities by region and pollutant over time.




Figure A.5: Baseline (pre-mitigation) emissions by region.




Figure A.6: Co-reduction relationships. The slope of each line represents the fraction of air pollutant emissions that are reduced given a fraction reduction in CO₂ emissions. Each dot is a single region-specific estimate embedded within a given Shared Socioeconomic Pathway (SSP). The co-reduction relationship for PM_{2.5} is taken as the average of black carbon and organic carbon.




Figure A.7: Radiative forcing. Radiative forcing over time in the business-as-usual scenario as well as the reference case and full RICE+AIR optima. Dotted lines sum to the solid purple line.




Figure A.8: Cost curves for air quality control. Cost curves by region and pollutant used in the simulation where RICE+AIR optimizes on both climate policy and air quality policy simultaneously. The curves terminate at the maximum technologically feasible level of reduction that can be achieved through end-of-pipe measures.

A.2 Supplementary Tables

Table A.1: Effect of valuing life-years differently in wellbeing terms in the optimization. Percent increase in optimal decarbonization rates in 2030 and 2050 and the difference in peak temperature between the reference case and the RICE+AIR case using the following approaches (see Note A.2 for details):

- The approach underlying our main results where the value of a life year is higher in poor compared to rich regions.
- Three approaches where life-year values are equal across regions.
- An approach where the value of a life-year is higher in rich compared to poor regions. Specifically, the value of a life increases with the square root of consumption.

	2030	2050	Difference in peak temperature ^ψ (°C)
Value richer less (main result)	52%	49%	0.39
Equal value 1 [*]	30%	33%	0.38
Equal value 2 ^φ	42%	45%	0.45
Equal value 3 [†]	32%	37%	0.33
Value richer more	17%	26%	0.27

^ψ Positive values indicate lower temperatures in the RICE+AIR case. Values are rounded.

^{*} Inequality aversion (η) = 1, pure rate of time preference = 1.5.

^φ Inequality aversion (η) = 1, pure rate of time preference = 2.

[†] The value of a life-year is $\alpha \times c^\eta$ (see Note A.2).

Table A.2: The modified objective function. Tables A.2a and A.2b show the percent change in results between the reference case and the full RICE+AIR case using two model versions that differ only in their objective function. “Model 1” is the model we implement whereas “Model 2” is identical except it uses the standard RICE2010 objective function. The difference in the two objective functions is described below the table. Unlike Tables A.2a and A.2b, which are within-model comparisons, Table A.2c compares results between the two models in the full RICE+AIR case of each. Together, the tables show that although our version of RICE recommends slightly faster decarbonization than the original version of the model, within both models the effect of accounting for health co-benefits and aerosol forcing is similar in magnitude.

Table A.2a: Within-model comparison: Optimal Decarbonization				
	2030	2050	2070	2090
Model 1	52%	49%	39%	22%
Model 2	48%	44%	35%	20%
Table A.2b: Within-model comparison: Temperature rise				
	2030	2050	2070	2090
Model 1	< 1%	-2%	-5%	-8%
Model 2	< 1%	-2%	-3%	-6%
Table A.2c: Between-model comparison: Percent change from Model 1 to Model 2 in the RICE+AIR case				
	2030	2050	2070	2090
Decarbonization Rate	-17%	-17%	-17%	-15%
Temperature	< -1%	1%	2%	5%

Note: the modified objective function explained

Consider the following general weighted objective function that is often used in climate-economy models:

$$W(c_{it}) = \sum_{it} \mathfrak{N}_{it} \frac{L_{it}}{(1 + \rho)^{10t}} \frac{c_{it}^{(1-\eta)}}{1 - \eta} . \quad (\text{A.1})$$

If the weights that appear in this equation, $\mathfrak{N}$, are proportional to the inverse of the marginal utility of consumption they are called Negishi weights. Negishi weights are used in many climate-economy models, including early versions of RICE, where they were introduced to impose constraints on capital flows (Nordhaus and Yang 1996). RICE2010 does not require weights for this purpose (because regions are autarkic), but they are still used so that the maximization as market simulation principle holds (Nordhaus and Sztorc 2013).

Our objective function (Eq. 2.1 in the Methods section) is identical to Eq. A.1 but without any weights ($\mathfrak{N} = 1$). This is because we do not suppose our results represent a market simulation since we do not expect the mitigation rates to emerge as the result of an unregulated market. Furthermore, Negishi weights distort time preferences (Dennig and Emmerling 2017) and the inter-regional trade-off (Stanton 2011) in ways that are opaque and difficult to justify, both descriptively and normatively. We have explained this change in more detail previously (Dennig et al. 2015).

Table A.3: Radiative forcing coefficients by region in 2050 and 2100.

	SO ₂	NO _x	Black Carbon	Organic Carbon
2050				
USA	-0.0121	-0.0052	0.0739	-0.0088
Europe	-0.0121	-0.0052	0.0739	-0.0088
Japan	-0.0121	-0.0052	0.0739	-0.0088
Russia	-0.0129	-0.0048	0.0732	-0.0082
Eurasia	-0.0129	-0.0048	0.0732	-0.0082
China	-0.0112	-0.0056	0.0746	-0.0095
India	-0.0112	-0.0056	0.0746	-0.0095
North Africa & Middle East	-0.0094	-0.0065	0.0760	-0.0109
Sub-Saharan Africa	-0.0094	-0.0065	0.0760	-0.0109
Latin America	-0.0067	-0.0077	0.0780	-0.0129
Other High Income	-0.0121	-0.0052	0.0739	-0.0088
Other Asia	-0.0112	-0.0056	0.0746	-0.0095
2100				
USA	-0.0130	-0.0062	0.0726	-0.0103
Europe	-0.0130	-0.0062	0.0726	-0.0103
Japan	-0.0130	-0.0062	0.0726	-0.0103
Russia	-0.0140	-0.0058	0.0718	-0.0097
Eurasia	-0.0140	-0.0058	0.0718	-0.0097
China	-0.0121	-0.0065	0.0733	-0.0109
India	-0.0121	-0.0065	0.0733	-0.0109
North Africa & Middle East	-0.0101	-0.0072	0.0748	-0.0122
Sub-Saharan Africa	-0.0101	-0.0072	0.0748	-0.0122
Latin America	-0.0071	-0.0082	0.0771	-0.0140
Other High Income	-0.0130	-0.0062	0.0726	-0.0103
Other Asia	-0.0121	-0.0065	0.0733	-0.0109

Units are in W/m²/Gg emissions.

A.3 Supplementary Notes

Note A.1: Co-optimizing air quality and climate policy

Our main results assume that autonomous air quality control proceeds approximately as expected. However, in addition to exogenously specifying such control, there is capability in AIR to cost-optimize it through end-of-pipe measures. In this case, RICE+AIR selects the optimal combination of both air quality and climate policies. Upper limit emission reductions and associated cost curves for each pollutant and region are derived and projected based on data from the GAINS model using the maximum technically feasible reduction scenarios (Fig. A.8). We deflate the technology costs in GAINS by 50% to estimate the total costs to the economy, which are generally lower than technology costs (Vrontisi et al. 2016). Note that climate policy can reduce air pollutant emissions below what is technically feasible through end-of-pipe technologies alone – for example by switching from a coal plant to solar photovoltaic – and therefore may still lead to co-benefits even in regions that implement all available end-of-pipe measures. We assume that end-of-pipe air pollution technologies do not affect CO₂ emissions.

Note A.2: Life-year monetization and valuation

Our main results assume that an additional life-year lived is proportional to regional per capita consumption. One possible concern with this approach, where the monetization method is based on regional per capita consumption, is that it may appear to imply that an additional life-year lived counts for less in a poorer region than in a richer region. However, our modeling avoids this worrisome implication because of the degree of concavity in the objective function in consumption (Eq. 2.1 in the Methods). This concavity, represented by the elasticity of marginal utility, η , results in diminishing marginal utility of consumption (inequality aversion) in the objective function. Specifically, impacts are valued differently depending on the consumption level, so that an impact of x dollars is counted in the objective as $x * c_i^{-\eta}$ when it affects an individual with consumption c_i . As a result, the wellbeing value of a life year (*WLY*) to a person with consumption c_i is computed as follows:

$$WLY = VOLY * c_i^{-\eta} = 2c_i^{(1-\eta)}. \quad (A.2)$$

To reiterate, the *VOLY* is what a life-year is worth in dollar terms whereas the *WLY* is the wellbeing value of the *VOLY* after it passes through the objective function. Eq. A.2 shows that when $\eta > 1$, each life-year is valued more highly in the objective function in poor regions than in rich regions. As a result, because the standard value of η we use is 1.5, our objective function actually assigns greater value to each life-year gained in poorer regions, despite each life-year having a lower monetized value in those regions.

With this in mind, we designed a further sensitivity test (Table A.1) to explore the implications of alternative approaches that each ensure that life-years are weighted exactly equally in rich and poor regions in the objective function, despite having different monetized *VOLY*s. The first is to set $\eta = 1$, which has the desired effect; however, since this method affects the objective function generally, it has the broader influence of increasing optimal decarbonization independent of the health impacts. Therefore, in the second approach we compensate for this issue by also adjusting the rate of time preference so the reference case with $\eta = 1$ approximates the standard reference case with $\eta = 1.5$. For the final approach, we run the model with the assumption that the value of a life-year is $\alpha * c^\eta$. By Eq. A.2 this leads to a social value of a life-year that is equal to α . In particular, it is identical across income groups. We choose α so that the total discounted social value of the health co-benefits are equal to what they are when $VOLY = 2c$ under the full RICE+AIR model optimum.

In contrast and in addition to these sensitivities, Table A.1 also reports results when rich lives are valued more than poor lives in the objective function by setting $VOLY = \alpha * c^2$. By Eq. A.2 this results in a social value of a life-year that is equal to $\alpha * \sqrt{c}$, which is increasing in income (i.e. the social value of a life increases with the square root of consumption). Here again we choose α so that the total discounted social value of the health co-benefits are equal to what they are when $VOLY = 2c$ under the full RICE+AIR model optimum.

Results of all sensitivity analyses described in this section show that the strong effect of adding health co-benefits remains under all these approaches, although it is generally somewhat reduced.

Note A.3: The FAIR climate model

In one sensitivity presented in the main text we performed an experiment where we ran RICE+AIR using the Finite Amplitude Impulse Response (FAIR) climate model (Millar et al. 2017), version 1.0, as an alternative to RICE's native climate model. However, these two models are not straightforwardly compatible, in large part because RICE runs on decadal time steps while FAIR is annual. Creating yearly versions of the socioeconomic, emission, and climate damage modules from RICE would address this issue, but such an exercise requires significant model development and validation efforts beyond the scope of this paper.

We instead created an iterative version of RICE+AIR with FAIR that simulates a direct coupling between the two models. This iterative approach proceeds in several steps:

1. For a given global carbon tax, use a deterministic version of RICE+AIR to calculate CO₂ emissions and aerosol radiative forcing values.
2. Interpolate the CO₂ and aerosol radiative forcing decadal values into annual values.
3. Use these interpolated results and the same exogenous radiative forcing scenario to calculate projected temperatures with FAIR (we import the aerosol forcing estimates from RICE+AIR because endogenous aerosol forcing is not represented in FAIR version 1.0).
4. Run RICE+AIR with each decadal value from the updated FAIR temperature projections, which changes the estimated climate damages, economic output, and CO₂ emission levels.
5. Use these updated CO₂ emission and aerosol forcing values from RICE+AIR to again run FAIR.

We repeat the process outlined above until the results converge and can thus optimize this iterative version of the model.

Note A.4: FUND+AIR

Like RICE/DICE, FUND is one of the three leading climate-economy models used by the US Interagency Working Group to estimate the social cost of carbon (Interagency Working Group on the Social Cost of Carbon 2016). (The third model, PAGE, is less widely used). FUND is a 16-region global model that, unlike RICE, was not designed to optimize. Instead, FUND is mainly used to either explore the impacts of climate policy scenarios or compute the stream of climate-related damages from a pulse emission of CO₂. FUND has different world regions, a different economic framework, a different climate module and a different specification of climate damages when compared to RICE. It thus provides an important opportunity to explore model uncertainty. The key features of FUND have been described elsewhere (Tol 1997; Waldhoff et al. 2014); here we limit the discussion to the modifications we made to optimize the model and link it to the AIR module.

Optimization: We optimize FUND through 2150 by finding the carbon tax pathway that maximizes Eq. 2.1 of the Methods, with the same basic trade-off as in RICE between mitigation costs and climate damages. Results were computed with the same discounting parameters used in the RICE+AIR model runs – an inequality aversion and time preference of 1.5 and 1.5%, respectively.

Adding the AIR module: We implemented FUND+AIR through the same approach described above for RICE+AIR, as summarized in two steps. First, we removed all exogenous aerosols and their precursors, and endogenized the (pre-mitigation) emissions (see Eqs. 2.2-2.4 in the Methods section). Second, we allowed emissions to change as a result of CO₂ mitigation (Eq. 2.6 in the Methods section). The difference in pre- and post-mitigation emissions determines the impact of decarbonization on radiative forcing and human health using the same methods described above for RICE+AIR (Eqs. 2.8-2.23). All relevant parameters were fit to the FUND regions, which differ somewhat from the RICE regions. For consistency with RICE+AIR, each life-year gained was valued at two years of per capita consumption.

APPENDIX B

Supplementary Information for Chapter 3

B.1 Supplementary Figures




Figure B.1: Additional probabilistic climate hindcasts. Hindcasts over the model calibration period (1850-2017) using the S-FAIR (purple), S-FUND (teal), S-Hector (red), and S-MAGICC (orange) climate models. **a-d**, The annual atmospheric carbon dioxide concentration, with different shapes identifying the Law Dome ice core (triangles) and Mauna Loa (circles) calibration data. **e-h**, Annual ocean carbon flux. **i-l**, Global ocean heat content. In all panels, yellow shapes represent observations. Solid black center line depicts the mean response across 100,000 model runs, with colored regions spanning the 95% predictive credible interval.




Figure B.2: Additional SC-CH₄ estimates under different model structures and treatments of parametric uncertainty. **a**, Distributions from the S-FAIR climate model for SC-CH₄ experiment using a constant 3% consumption discount rate under RCP8.5. Different distribution line types identify the FUND (solid) and DICE (dashed) IAMs, with colored shapes along the x-axis marking the estimated means of the different distributions. Colors identify each SC-CH₄ experiment, with our main specification using Bayesian calibration (“Main SC-CH₄ estimate”, red), outdated radiative forcing equations that disregard methane’s shortwave absorption (“Outdated CH₄ radiative forcing”, orange), neglecting posterior relationships and sampling each parameter independently (“Remove parameter relationships”, green), and sampling the ECS distribution used for U.S. SC-CH₄ estimates while fixing other uncertain climate model parameters at their mean posterior values (“U.S. climate sensitivity”, blue). **b**, same as **a** but for the S-FUND climate model. **c**, same as **a** but for the S-Hector climate model.




Figure B.3: Using the RCP2.6 scenario keeps expected temperature projections below 2 °C and reduces SC-CH₄ estimates relative to RCP8.5. a-d, Modeled annual average global surface temperature anomalies relative to the 1861-1880 mean under RCP2.6. Red dashed vertical lines identify the end of the calibration period (1850-2017). Yellow circles represent temperature observations and solid black center line depicts the mean response across 100,000 model runs. Colored regions span the 95% predictive credible intervals and horizontal gray lines identify the UN Paris Agreement’s 1.5 °C and 2 °C global temperature targets. **e,** SC-CH₄ distributions for FUND (solid line) and DICE (dashed line) using a constant 3% consumption discount rate under RCP2.6. White circles and diamonds show FUND and DICE’s expected SC-CH₄ for the two RCP scenarios after pooling together results from the four climate models. The percentage value identifies the percent change in each IAM’s expected SC-CH₄ that occurs when switching from RCP8.5 to RCP2.6, with the arrow identifying the direction of the change. In all panels, different colors identify the S-FAIR (purple), S-FUND (teal), S-Hector (red), and S-MAGICC (orange) models.




Figure B.4: Robustness of SC-CH₄ distributions to wider prior assumptions, and posterior methane cycle parameter relationships. **a**, SC-CH₄ distributions for FUND (solid line) and DICE (dashed line) depicting the main results (red) and a sensitivity analysis that uses wider prior parameter distributions during the model calibrations (blue). The SC-CH₄ distributions correspond to a constant 3% consumption discount rate under RCP8.5 and pool each IAM's SC-CH₄ estimates across the four climate models. Colored circles show the estimated multi-model mean SC-CH₄ for the two scenarios. **b**, Posterior parameter relationship between natural methane emission rates and methane's initial tropospheric lifetime for the S-MAGICC climate model. **c**, Posterior relationships between the uncertain methane cycle parameters depicted in **b** with S-MAGICC's estimated SC-CH₄ for RCP8.5 under a constant 3% consumption discount rate. Different sized diamonds and circles identify methane's initial tropospheric lifetime for DICE (blue) and FUND (red). Both **b** and **c** depict 5,000 randomly selected posterior estimates, with loess-smoothed curves (white lines) helping illustrate the relationship between x and y-axis values.




Figure B.5 | Additional models showing non-linear climate parameter relationships constrain SC-CH₄ estimates. Parameter and SC-CH₄ relationships for the S-FAIR (top row), S-FUND (middle row), and S-Hector (bottom row) climate models. **a-c**, Posterior relationships between four uncertain climate parameters. Different sized circles correspond to different terrestrial carbon pool respiration-temperature sensitivity values (with higher values signaling increasing heterotrophic respiration of carbon dioxide with temperature). Each dot's color scales with the vertical rate of heat diffusion into the ocean. **d-f**, Posterior relationships between three of the uncertain climate parameters depicted in **a-c** with the SC-CH₄ estimated for RCP8.5 under a constant 3% consumption discount rate. Different sized diamonds and circles identify respiration temperature sensitivity values for DICE and FUND. Each point's color scales with the aerosol radiative forcing factor (with higher values signaling stronger aerosol cooling). Each panel depicts 5,000 randomly selected posterior estimates, with loess-smoothed curves (white lines) helping to illustrate the relationship between x and y-axis values.




Figure B.6: Neglecting relationships between posterior parameter estimates increases probabilistic temperature projection uncertainty. a-c, Modeled annual average global surface temperature anomalies relative to the 1861-1880 mean using the S-FAIR (purple), S-FUND (teal), and S-Hector (red) climate models. Yellow circles represent temperature observations and solid black center line depicts the mean baseline response across 100,000 model runs. Colored regions outlined by dashed lines span the 95% predictive credible interval. The outer gray colored regions bound the 95% predictive credible intervals for model projections that remove parameter relationships by sampling each marginal posterior distribution independently. The inset figures depict estimated temperature distributions for 2050 and 2100 using the full set of model runs from the baseline (colored) and no parameter relationship (gray) scenarios.




Figure B.7: Strong positive relationship between the SC-CH₄ and ECS under the U.S. SC-CH₄ estimation framework. **a**, Solid colored lines depict posterior ECS distributions for the S-FAIR (purple), S-FUND (teal), S-Hector (red), and S-MAGICC (orange) climate models. The dark gray dashed line shows the ECS distribution used for official U.S. SC-CH₄ estimates. Colored star shapes along the x-axis identify the estimated mean for the different distributions. **b-e**, SC-CH₄ values estimated for a constant 3% consumption discount rate under RCP8.5 following the U.S. SC-CH₄ estimation framework's treatment of parametric climate uncertainty. ECS values are sampled from the same distribution used to estimate the U.S. SC-CH₄ values. Other uncertain climate model and statistical process parameters remain fixed at their posterior mean values. Each panel depicts 2,000 randomly selected SC-CH₄ estimates for DICE (circles) and FUND (diamonds), with loess-smoothed curves (white lines) helping illustrate the relationship between the ECS and SC-CH₄.

B.2 Supplementary Tables

Table B.1: Past SC-CH₄ estimates and their treatment of parametric climate uncertainty

Previous Study	SC-CH ₄ Estimate (95 th percentile)	Discounting Framework	Emission Year	IAMs Used	Treatment of Parametric Climate Uncertainties
Hope (2005)	\$175 (N/A)	Ramsey	2000	PAGE95	250 Latin hypercube samples of more than 70 uncertain economic and climate parameters. Uncertain parameters represented by pre-specified, independent triangular distributions.
Hope (2006)	\$126 (\$315)	Ramsey	2001	PAGE2002	Latin hypercube sampling over 31 uncertain economic and climate parameters represented by pre-specified, independent triangular distributions.
Marten and Newbold (2012)	\$1100 (N/A)	Constant (3%)	2020	DICE2007 coupled to MAGICC climate model	Monte Carlo simulation using Roe and Baker equilibrium climate sensitivity (ECS) distribution consistent with IPCC AR4 Report statements about ECS uncertainty, truncated between 0-10 °C. Other climate parameters held constant.
Waldhoff et al. (2014)	\$313 (N/A)	Ramsey	2010-2019	FUND v3.9	Point estimates using different ECS values (2.0, 3.0, and 4.5 °C) and with terrestrial carbon sink temperature feedbacks switched off. Other climate parameters held constant.
Marten et al. (2015)	\$1200 (\$3200)	Constant (3%)	2020	FUND (unmodified) plus DICE and PAGE coupled to a one-box methane cycle model with a constant atmospheric decay rate.	Default uncertain climate parameters in FUND and PAGE sampled independently. Each IAM also sampled a Roe and Baker ECS distribution consistent with IPCC AR4 Report statements about ECS uncertainty, truncated between 0-10 °C. Other climate parameters held constant.

This table has been adapted from Table 1 in Marten et al. (2015) to highlight SC-CH₄ estimates from previous studies that also account for parametric climate uncertainties. All SC-CH₄ values have been converted to 2007 U.S. dollars. The values in the second column show a study's central SC-CH₄ estimate, with lower values in parenthesis corresponding to the 95th percentile estimate when reported.

Table B.2: Uncertain climate model and statistical process parameters with their prior distributions

Uncertain Model Parameter	Units	Prior Distributions (main results)	Wider Prior Distributions (sensitivity analysis)
Equilibrium climate sensitivity	°C	Cauchy(3,2) truncated between 0 and 10	Cauchy(3, 2.4) truncated between 0 and 12
Ocean vertical heat diffusivity	cm ² /s	Log-normal(1.1, 0.3)	Log-normal(1.1, 0.36)
Aerosol radiative forcing scaling factor	—	Triangular(0, 3, 1)	Triangular(0, 3.6, 1)
CH ₄ direct radiative forcing scaling factor	—	Uniform(0.72, 1.28)	Uniform(0.576, 1.536)
Forcing increase from CO ₂ doubling*	W/m ²	Uniform(2.97, 4.45)	Uniform(2.376, 5.34)
Ocean carbon diffusivity	m/yr	Uniform(0, 200)	Uniform(0, 240)
Respiration temperature sensitivity	—	Uniform(1, 5)	Uniform(1, 6)
Carbon fertilization factor	—	Uniform(0, 1)	Uniform(0, 1.2)
Natural CH ₄ emissions [†]	Mt/yr	Uniform(150, 484)	Uniform(120, 580.8)
CH ₄ soil sink lifetime [‡]	years	Uniform(96, 144)	Uniform(76.8, 172.8)
CH ₄ stratospheric sink lifetime [‡]	years	Uniform(100, 200)	Uniform(80, 240)
Constant tropospheric CH ₄ lifetime [§]	years	Uniform(6.7, 13.5)	Uniform(5.36, 16.2)
Initial tropospheric CH ₄ lifetime [‡]	years	Uniform(6.7, 13.5)	Uniform(5.36, 16.2)
Initial global mean surface temperature	°C	Normal(0, 1)	Normal(0, 1.2)
Initial ocean heat uptake	10 ²² J	Uniform(-100, 0)	Uniform(-120, 0)
Initial atmospheric CO ₂ concentration	ppm	Uniform(275, 281)	Uniform(220, 337.2)
Initial atmospheric CH ₄ concentration	ppb	Uniform(691, 753)	Uniform(552.8, 903.6)
Initial atmospheric N ₂ O concentration	ppb	Uniform(264, 282)	Uniform(211.2, 338.4)
AR(1) standard deviation (temperature)	°C	Uniform(0, 0.2)	Uniform(0, 0.24)
AR(1) standard deviation (ocean heat content)	10 ²² J	Uniform(0, 4)	Uniform(0, 4.8)
AR(1) Autocorrelation (temperature)	—	Uniform(-0.99, 0.99)	Uniform(-0.99, 0.99)
AR(1) Autocorrelation (ocean heat content)	—	Uniform(-0.99, 0.99)	Uniform(-0.99, 0.99)
Measure of CAR(1) correlation memory (CO ₂)	—	Uniform(0.01, 11.5)	Uniform(0.008, 13.8)
Measure of CAR(1) correlation memory (CH ₄)	—	Uniform(0.01, 11.5)	Uniform(0.008, 13.8)
Continuous white noise process variance (CO ₂)	—	Uniform(0, 200)	Uniform(0, 240)
Continuous white noise process variance (CH ₄)	—	Uniform(0, 800)	Uniform(0, 960)

* We follow Smith et al. (2018) and scale CO₂ radiative forcing in all time periods to be consistent with the sampled forcing increase from CO₂ doubling.

† Uncertain methane cycle parameter in all climate models except S-FUND.

‡ Uncertain methane cycle parameter in S-Hector and S-MAGICC only.

§ Uncertain methane cycle parameter in S-FAIR and S-FUND only.

Descriptions of the uncertain climate model and statistical process parameters. The third column (“Prior Distributions”) lists the prior parameter distributions used for the baseline model calibrations. The fourth column (“Wider Prior Distributions”) corresponds to an analysis exploring the SC-CH₄’s sensitivity to model calibrations based on wider and more diffuse prior parameter distributions.

Table B.3: Model and scenario specific SC-CH4 estimates

IAM	Climate Model	RCP8.5					RCP2.6	Neglect Parameter Relationships	Outdated Methane Forcing	U.S. Climate Sensitivity
		Constant 2.5%	Constant 3%	Constant 5%	Ramsey 1.5%, 1.0	Ramsey 1.5%, 1.5				
DICE	S-FAIR	1760 (852-3070)	1390 (712-2350)	745 (432-1150)	1220 (618-2090)	855 (466-1400)	1000 (553-1600)	1490 (583-2950)	1200 (623-2010)	1720 (540-4190)
	S-FUND	1450 (715-2530)	1160 (603-1950)	635 (376-974)	1020 (526-1740)	723 (403-1170)	893 (500-1430)	1240 (487-2440)	978 (514-1630)	1400 (425-3480)
	S-Hector	1750 (842-3110)	1390 (705-2380)	736 (424-1150)	1210 (609-2120)	847 (460-1410)	1030 (569-1700)	1480 (572-2970)	1160 (597-1980)	1700 (537-4150)
	S-MAGICC	1740 (907-2870)	1390 (756-2220)	747 (452-1110)	1211 (653-1960)	852 (490-1330)	1050 (608-1600)	1470 (623-2770)	1180 (651-1880)	1650 (550-3870)
FUND	S-FAIR	620 (281-1050)	535 (254-886)	340 (171-541)	548 (255-915)	454 (220-744)	408 (161-704)	566 (173-1120)	452 (213-745)	630 (151-1460)
	S-FUND	517 (216-881)	449 (202-748)	292 (141-468)	459 (200-773)	384 (178-632)	368 (136-642)	473 (132-946)	371 (165-619)	518 (102-1230)
	S-Hector	617 (274-1050)	531 (247-891)	334 (164-539)	544 (248-920)	450 (213-747)	421 (163-737)	561 (171-1120)	436 (201-732)	625 (153-1440)
	S-MAGICC	619 (290-1010)	534 (263-856)	339 (175-529)	547 (264-884)	454 (227-722)	423 (173-711)	562 (184-1070)	446 (216-712)	606 (150-1370)
Column Average		1134 (547-1946)	922 (468-1535)	521 (292-808)	845 (422-1425)	627 (332-1019)	699 (358-1141)	980 (366-1923)	778 (398-1289)	1106 (326-2649)

Reported values show the average SC-CH4 estimate for a specific scenario-model combination, with lower values in parenthesis showing the corresponding 5-95% predictive credible interval . The five columns under the RCP8.5 heading depict estimates under constant consumption discounting (2.5%, 3%, and 5%) and Ramsey discounting (pure rate of time preference = 1.5%, elasticity of marginal utility of consumption = 1.0 and 1.5). The remaining columns show alternative SC-CH4 estimates using a constant 3% consumption discount rate for the RCP2.6 scenario (“RCP2.6”), neglecting posterior relationships and sampling each parameter independently (“Neglect Parameter Relationships”), outdated radiative forcing equations that disregard methane’s shortwave absorption (“Outdated Methane Forcing”), and sampling the ECS distribution used for U.S. SC-CH4 estimates while fixing other uncertain parameters at their mean posterior values (“U.S. Climate Sensitivity”). The bottom row shows each column’s average SC-CH4 estimate.

APPENDIX C

Supplementary Information for Chapter 4

C.1 Supplementary Figures




Figure C.1 | The DICE, FUND, and PAGE climate models exhibit limited temperature and sea level hindcast skill under their default parameterizations. a-c, Predicted global mean surface temperature anomalies relative to the 1861-1880 mean using the DICE (**a**, light orange), FUND (**b**, green), and PAGE climate models. Center lines outlined in black depict the mean model response for 10,000 runs that sample each model’s default, uncorrelated climate parameter distributions and the equilibrium climate sensitivity distribution used in official U.S. SCC estimates. Colored regions span the 95% confidence interval. The solid red line depicts temperature observations with lighter shading for the 2σ observation error range. **d-f**, The same as in **a-c**, but for global mean sea level relative to the 1961-1990 mean.




Figure C.2 | Posterior climate model parameter distributions. **a**, Posterior equilibrium climate sensitivity distributions for the DICE-Climate (orange), FUND-Climate (green), PAGE-Climate (blue), and SNEASY+BRICK (black) models under the baseline Bayesian calibration framework in addition to a scenario where SNEASY+BRICK is recalibrated without sea level observations (black). The equilibrium climate sensitivity distribution used for official U.S. SCC estimates is shown in red. Vertical lines identify each distribution's 95th percentile equilibrium climate sensitivity value, with the gray dashed line showing the 95th percentile estimate from Sherwood et al. (2020). **b-c**, Two examples of default versus posterior parameter distributions for the FUND climate model. Uniform prior distributions used in the model calibration shown as black dashed rectangles. **d-e**, Same as in **b-c** but for the PAGE climate model.




Figure C.3 | Posterior parameter correlations emerge after calibrating the DICE-Climate, FUND-Climate, and PAGE-Climate models to past observations. a-b, Two examples of posterior parameter correlations for DICE-Climate. **c-d**, Same as in **a-b** but for FUND-Climate. **e-f**, Same as in **a-b** but for the PAGE-Climate.




Figure C.4 | Climate model structure and parameter correlations influence out-of-sample sea level hindcast skill. Predicted global mean sea level relative to the 1961-1990 mean using the DICE-Climate (a, light orange), FUND-Climate (b, green), PAGE-Climate (c, blue), and SNEASY+BRICK (d, gold) models. Each model is calibrated to observations covering the period 1850-1970 and then performs an out-of-sample hindcast test from 1970-2017 (gray shaded region). Center lines outlined in black depict the mean model response for 10,000 posterior parameter samples and the edges of the colored regions span the 95% predictive credible intervals. The dashed black center line shows the mean model response when calibrating each model to the full set of observations from 1850-2017. Purple lines show the 95% predictive credible intervals for a scenario that neglects posterior parameter correlations by sampling each calibrated parameter distribution independently. The solid red line depicts global mean sea level observations with lighter shading for the 2σ observation error range.




Figure C.5 | Forcing scenarios and different treatments of parametric uncertainty influence a variety of model responses to a 100 GtC emissions pulse. Different climate model impulse response distributions arising from a 100 GtC emission pulse in 2020 using the DICE-Climat (top row), FUND-Climat (second row), PAGE-Climat (third row), and SNEASY+BRICK (bottom row) models. Each distribution is based on 10,000 model runs. **a-d**, Peak temperature increase between 2020-2300 from a 100 GtC emission pulse for RCP8.5 (green), RCP2.6 (blue), and a scenario that neglects posterior parameter correlations by sampling the marginal posterior distributions independently (purple). For DICE-Climat, FUND-Climat, and PAGE-Climat, the red lines depict a scenario sampling the models' default, uncorrelated parameter distributions and the equilibrium climate sensitivity distribution used for official U.S. SCC estimates. Under this scenario, the DICE-Climat and PAGE-Climat run on ten-year time steps. For SNEASY+BRICK, the red lines depict a similar scenario that neglects posterior correlations and also samples the U.S. equilibrium climate sensitivity distribution. **e-h**, same as in **a-d** but depicting the number of years required to achieve the maximum temperature increase. Due to their ten-year time steps, DICE-Climat and PAGE-Climat report results from the "U.S. Uncertainty" scenario as a histogram normalized to the distribution density values. **i-l**, same as in **a-d** but for the 2120 sea level increase occurring one hundred years after the emissions pulse.




Figure C.6 | Climate model parameter correlations reduce uncertainty in the temperature response to cumulative carbon dioxide emissions. Predicted global mean surface temperature increase relative to cumulative carbon dioxide emissions using the DICE-Climate (**a**, light orange), FUND-Climate (**b**, green), PAGE-Climate (**c**, blue), and SNEASY+BRICK (**d**, gold) models. The lines corresponding to each model's specific color depict individual runs from a 1,000 member ensemble. For each posterior parameter sample, the model is forced with a unique emissions scenario based on an inversion calculation such that atmospheric carbon dioxide concentrations increase 1% per year when all other external forcings remain fixed at their 1850 values. The center white line outlined in black shows the mean temperature response to cumulative emissions for the ensemble. Purple lines similarly show the individual runs from a 1,000 member ensemble that neglects the posterior parameter correlations by sampling the marginal posterior distributions independently, with the dashed center black line showing the mean response.




Figure C.7 | Sea level response dynamics and uncertainty levels resulting from a 100 GtC emission pulse vary with climate model parameter choice. Predicted global mean sea level from a 100 GtC emission pulse in 2020 under the RCP8.5 scenario using the DICE-Climate (a, orange), FUND-Climate (b, green), PAGE-Climate (c, blue), and SNEASY+BRICK (d, gold) models. The lines corresponding to each model's specific color depict 1,000 randomly sampled runs from a 10,000 member ensemble using the posterior parameter samples, with the center white line outlined in black showing the mean impulse response for the full ensemble. Purple lines similarly show 1,000 randomly sampled runs from a 10,000 member ensemble that neglects the posterior parameter correlations by sampling the marginal posterior distributions independently. For DICE-Climate, FUND-Climate, and PAGE-Climate, the solid red line shows the mean model response from 10,000 runs sampling the models' default, uncorrelated parameter distributions and the equilibrium climate sensitivity distribution used for official U.S. SCC estimates. Under this scenario, DICE-Climate and PAGE-Climate run on ten-year time steps. Dashed red lines show the 95% confidence interval. For SNEASY+BRICK, the red lines depict a similar scenario that neglects posterior correlations and also samples the U.S. equilibrium climate sensitivity distribution.




Figure C.8 | Social cost of carbon distributions under different discount rates. SCC distributions estimated with the SB-DICE (a), SB-FUND (b) and SB-PAGE (c) integrated assessment models under the RCP8.5 scenario, with each distribution corresponding to 10,000 SCC estimates for an emission pulse in 2020. Solid-lined distributions use constant consumption discount rates of 2.5% (red), 3% (orange), 5% (green), and 7% (purple). The dashed-line distributions depict SCC estimates under a Ramsey discounting framework with a 1.5% pure rate of time preference and a consumption elasticity of marginal utility of 1.0 (pink) and 1.5 (blue).




Figure C.9 | Temperature impulse response dynamics from a 100 GtC emission pulse under the RCP2.6 scenario. Predicted global mean surface temperature increase from a 100 GtC emission pulse in 2020 under the RCP2.6 scenario using the DICE-Climate (a, light orange), FUND-Climate (b, green), PAGE-Climate (c, blue), and SNEASY+BRICK (d, gold) models. The lines corresponding to each model's specific color depict 1,000 randomly sampled runs from a 10,000 member ensemble using the posterior parameter samples, with the center white line outlined in black showing the mean impulse response for the full ensemble. Purple lines similarly show 1,000 randomly sampled runs from a 10,000 member ensemble that neglects the posterior parameter correlations by sampling the marginal posterior distributions independently. For DICE-Climate, FUND-Climate, and PAGE-Climate, the solid red line shows the mean model response from 10,000 runs sampling the models' default, uncorrelated parameter distributions and the equilibrium climate sensitivity distribution used for official U.S. SCC estimates. Under this scenario, DICE-Climate and PAGE-Climate run on ten-year time steps. Dashed red lines show the 95% confidence interval. For SNEASY+BRICK, the red lines depict a similar scenario that neglects posterior correlations and also samples the U.S. equilibrium climate sensitivity distribution.




Figure C.10 | Climate model projections and social cost of carbon distributions differ markedly between the RCP2.6 and RCP8.5 scenarios. **a**, Global mean surface temperature anomaly projections (relative to the 1861-1880 mean) for RCP8.5 (red) and RCP2.6 (blue) scenarios using the SNEASY+BRICK model. Center lines outlined in black depict the mean model response for 10,000 posterior parameter samples and the edges of the colored region outlined in dashed black lines span the 95% predictive credible intervals. The dashed red line shows the mean projection under RCP8.5 for a scenario where SNEASY+BRICK is recalibrated without sea level observations. The middle-shaded regions span the 95% interval for climate projections that neglect the posterior correlations by independently sampling each marginal posterior distribution. The outermost, lightest shaded regions show the 95% interval for projections that neglect posterior parameter correlations and sample the equilibrium climate sensitivity distribution specified for official U.S. SCC estimates. The white solid line in the calibration period depicts the temperature observations. **b**, The same as in **a**, but for global mean sea level relative to the 1961-1990 mean. **c-e**, SCC estimates for RCP2.6 and RCP8.5 under the baseline Bayesian calibration (solid lines) and scenario that removes parameter correlations (dashed lines) for SB-DICE (**a**), SB-FUND (**b**), and SB-PAGE (**c**). Vertical lines show each distribution's 95th percentile SCC estimate.




Figure C.11 | Equity-weighted social cost of carbon distributions for different geographic regions and treatments of parametric climate uncertainties. **a-c**, Equity-weighted SCC estimates for the United States using the SB-RICE, SB-FUND, and SB-PAGE integrated assessment models under the RCP8.5 scenario. Solid-lined distributions use 1.5 for the consumption elasticity of marginal utility while dashed-line distributions use a value of 1.0 (with higher values representing increasing preferences for equality in consumption and risk aversion). All estimates use a 1.5% pure rate of time preference. Distributions with each model’s specific color depict baseline SCC estimates using SNEASY+BRICK’s calibrated parameter samples (“Baseline Estimate”). Purple distributions show SCC estimates based on climate projections that neglect the posterior correlations found in the “Baseline Estimate” scenario (“Remove Correlations”). Red lines depict SCC estimates for a scenario that neglects posterior parameter correlations and samples the equilibrium climate sensitivity distribution specified for official U.S. SCC estimates (“U.S. Uncertainty”). **d-f**, The same as in **a-c** but for Africa. **d-f**, The same as in **a-c** but for Latin America.

C.1 Supplementary Tables

Table C.1 | Uncertain parameter distributions for DICE-Climate’s default settings and the prior distributions used for the Bayesian calibration.

Uncertain Model Parameters	Units	Default Model Distributions	MCMC Prior Distributions
Equilibrium climate sensitivity	°C	Distribution from official U.S. SCC estimates *	Cauchy(3,2) truncated between 0 and 10
Aerosol radiative forcing scaling factor	—	—	Triangular(0, 1, 3)
Atmospheric temperature speed of adjustment coefficient	—	—	Uniform(0, 1)
Heat loss from atmosphere to oceans coefficient	—	—	Uniform(0, 1)
Heat gain by deep oceans coefficient	—	—	Uniform(0, 1)
Carbon cycle transition coefficient (atmosphere to shallow ocean)	%	—	Uniform(0, 0.99)
Carbon cycle transition coefficient (shallow ocean to deep ocean)	%	—	Uniform(0, 0.99)
Equilibrium CO ₂ concentration in atmosphere.	GtC	—	Uniform(580, 596)
Equilibrium CO ₂ concentration in biosphere/shallow oceans.	GtC	—	Uniform(1350, 1650)
Equilibrium CO ₂ concentration in deep oceans.	GtC	—	Uniform(9000, 11000)
Thermal expansion adjustment rate	m °C ⁻¹ year ⁻¹	—	Uniform(0, 0.0094)
Glaciers and small ice caps melt rate	m °C ⁻¹ year ⁻¹	—	Uniform(0, 0.0377)
Greenland ice sheet melt rate	m °C ⁻¹ year ⁻¹	—	Uniform(0, 0.0021)
Antarctic ice sheet melt rate	m °C ⁻¹ year ⁻¹	—	Uniform(0, 0.00012)
Antarctic ice sheet temperature threshold	°C	—	Uniform(0, 3.3)
Glaciers and small ice caps asymptotic sea level rise	m	—	Uniform(0.31, 0.53)
Greenland ice sheet asymptotic sea level rise	m	—	Uniform(7.16, 7.56)
Antarctic ice sheet asymptotic sea level rise	m	—	Uniform(57, 58.3)
Initial global mean surface temperature anomaly	°C	—	Normal(0, 1) truncated between -0.3 and 0.3
Initial lower ocean temperature anomaly	°C	—	Normal(0, 1) truncated between -0.3 and 0.3
Initial CO ₂ concentration in atmosphere.	GtC	—	Uniform(599, 614)
Initial CO ₂ concentration in biosphere/shallow oceans.	GtC	—	Uniform(0, 6.4)
Increase in atmospheric CO ₂ concentration between initial condition and period two ‡	GtC	—	Uniform(1359, 1661)
Initial CO ₂ concentration in deep oceans.	GtC	—	Uniform(9001, 11001)
Initial sea level rise from thermal expansion	m	—	Uniform(-0.046, 0.074)
Initial sea level rise from glaciers and small ice caps	m	—	Uniform(-0.0536, 0.0791)
Initial sea level rise from Greenland ice sheet	m	—	Uniform(-0.0086, 0.0144)
Initial sea level rise from Antarctic ice sheet	m	—	Uniform(-0.0476, 0.0559)
AR(1) standard deviation (temperature)	°C	—	Uniform(0, 0.2)
AR(1) standard deviation (sea level rise)	m	—	Uniform(0, 0.05)
AR(1) autocorrelation (temperature)	—	—	Uniform(-0.99, 0.99)
AR(1) autocorrelation (sea level rise)	—	—	Uniform(-0.99, 0.99)
CAR(1) white noise process variance (CO ₂)	—	—	Uniform(0, 200)
CAR(1) measure of correlation memory (CO ₂)	—	—	Uniform(0.01, 11.5)

* We sample $\frac{1.2}{1-N(0.62,0.034)}$ for the equilibrium climate sensitivity distribution used for official U.S. SCC estimates.

‡ CO₂ concentrations are fixed in DICE’s first two periods. We model uncertainty in the 1851 concentration as the increase over 1850.

Descriptions of the uncertain climate model and statistical process parameters. Since DICE-Climate does not have default uncertain parameter distributions, the third column (“Default Model Distributions”) only lists the equilibrium climate sensitivity distribution corresponding to the standard distribution all models sample for official U.S. SCC estimates. The fourth column (“MCMC Prior Distributions”) shows the prior parameter distributions for the Bayesian model calibrations using MCMC.

Table C.2 | Uncertain parameter distributions for FUND-Climate’s default settings and the prior distributions used for the Bayesian calibration.

Uncertain Model Parameters	Units	Default Model Distributions	MCMC Prior Distributions
Terrestrial biosphere temperature sensitivity	GtC °C ⁻¹	Gamma(4.92, 662.83)	Uniform(32, 20170)
Decay time in carbon cycle box 1	years	—	Uniform(0.9, 1.1)
Decay time in carbon cycle box 2	years	Normal(363, 90.75) bounded below at 0	Uniform(0, 884.5)
Decay time in carbon cycle box 3	years	Normal(74, 18.5) bounded below at 0	Uniform(0, 172.8)
Decay time in carbon cycle box 4	years	Normal(17, 4.25) bounded below at 0	Uniform(0, 41.5)
Decay time in carbon cycle box 5	years	Normal(2, 0.5) bounded below at 0	Uniform(0, 5.1)
Constant term in temperature e-folding time calculation	—	Normal(-42.7, 0.12)	Uniform(-45.6, -39.9)
Linear term in temperature e-folding time calculation	—	Normal(29, 0.026)	Uniform(27.5, 30.7)
Quadratic term in temperature e-folding time calculation	—	Normal(0.0015, 0.0016)	Uniform(-0.0062, 0.01)
Equilibrium climate sensitivity	°C	Distribution from official U.S. SCC estimates*	Cauchy(3,2) truncated between 0 and 10
Sea level rise e-folding time	years	Triangular(250, 500, 1000)	Uniform(237.5, 1050)
Sea level rise temperature sensitivity coefficient	—	Gamma(6, 0.4)	Uniform(0.1, 11.8)
Fraction of CO ₂ concentration in carbon cycle box 1 ^φ	—	—	Uniform(0.117, 0.143)
Fraction of CO ₂ concentration in carbon cycle box 2 ^φ	—	—	Uniform(0.18, 0.22)
Fraction of CO ₂ concentration in carbon cycle box 3 ^φ	—	—	Uniform(0.288, 0.352)
Fraction of CO ₂ concentration in carbon cycle box 4 ^φ	—	—	Uniform(0.225, 0.275)
Fraction of CO ₂ concentration in carbon cycle box 5 ^φ	—	—	Uniform(0.09, 0.11)
Aerosol radiative forcing scaling factor	—	—	Triangular(0, 1, 3)
Initial atmospheric CO ₂ concentration	ppm	—	Uniform(282, 288)
Share of 1850 atmospheric CO ₂ concentration in carbon cycle boxes 1-5 [‡]	—	—	Uniform(0, 0.99)
Initial CO ₂ stock in the terrestrial biosphere	GtC	—	Uniform(1.71×10^6 , 2.09×10^6)
Pre-industrial atmospheric CO ₂ concentration	ppm	—	Uniform(275, 281)
Initial global mean surface temperature anomaly	°C	—	Normal(0, 1) truncated between -0.3 and 0.3
Initial global mean sea level rise	m	—	Normal(0, 0.0242)
AR(1) standard deviation (temperature)	°C	—	Uniform(0, 0.2)
AR(1) standard deviation (sea level rise)	m	—	Uniform(0, 0.05)
AR(1) autocorrelation (temperature)	—	—	Uniform(-0.99, 0.99)
AR(1) autocorrelation (sea level rise)	—	—	Uniform(-0.99, 0.99)
CAR(1) white noise process variance (CO ₂)	—	—	Uniform(0, 200)
CAR(1) measure of correlation memory (CO ₂)	—	—	Uniform(0.01, 11.5)

* We sample $\frac{1.2}{1-N(0.62, 0.034)}$ for the equilibrium climate sensitivity distribution used for official U.S. SCC estimates.

φ The fraction of CO₂ concentrations in each carbon cycle box are normalized to sum to 1.

‡ The share of 1850 initial condition CO₂ concentrations in each carbon cycle box are normalized to sum to 1.

Descriptions of the uncertain climate model and statistical process parameters. The third column (“Default Model Distributions”) lists FUND-Climate’s original parameter distributions that would typically be sampled when estimating the SCC. In this column, the equilibrium climate sensitivity distribution is not specific to FUND-Climate but corresponds to the standard distribution all models sample for official U.S. SCC estimates. The fourth column (“MCMC Prior Distributions”) shows the prior parameter distributions for the Bayesian model calibrations using MCMC.

Table C.3 | Uncertain parameter distributions for PAGE-Climate’s default settings and the prior distributions used for the Bayesian calibration.

Uncertain Model Parameters	Units	Default Model Distributions	MCMC Prior Distributions
Fraction of CO ₂ emissions remaining in atmosphere	—	Triangular(0.25, 0.3, 0.35)	Uniform(0.225, 0.385)
Half-life of CO ₂ atmospheric residence	years	Triangular(50, 70, 100)	Uniform(45, 110)
Percent of CO ₂ emitted to the atmosphere	%	Triangular(57, 62, 67)	Uniform(51.3, 73.7)
CO ₂ concentration temperature sensitivity feedback	% °C ⁻¹	Triangular(4, 10, 15)	Uniform(3.6, 16.5)
CO ₂ concentration temperature sensitivity feedback limit	%	Triangular(30, 50, 80)	Uniform(27, 88)
Feedback response time to changes in radiative forcing	years	Triangular(10, 35, 65)	Uniform(9, 71.5)
Poles excess temperature change over equator	—	Triangular(1, 1.5, 2)	Uniform(0.9, 2.2)
Land excess temperature ratio to ocean	—	Triangular(1.2, 1.4, 1.6)	Uniform(1.08, 1.54)
Sea level rise with temperature	m °C ⁻¹	Triangular(0.7, 1.5, 3)	Uniform(0.63, 3.3)
Sea level rise asymptote	m	Triangular(0.5, 1, 1.5)	Uniform(0.45, 1.65)
Half-life of sea level rise	Years	Triangular(500, 1000, 1500)	Uniform(450, 1650)
Sea level rise in base year (1849)	m	Triangular(-0.048, 0.002, 0.052)	Normal(0, 0.0242)
Sulphate direct (linear) forcing effect in base year (1849)	Wm ⁻²	Triangular(-0.1, -0.004, 0)	Uniform(-0.11, 0)
Sulphate indirect (log) forcing effect for a doubling	Wm ⁻²	Triangular(-0.4, -0.06, 0)	Uniform(-0.44, 0)
Equilibrium climate sensitivity	°C	Distribution from official U.S. SCC estimates *	Cauchy(3,2) truncated between 0 and 10
Initial atmospheric CO ₂ concentration	ppb	—	Uniform(282000, 288000)
Pre-industrial atmospheric CO ₂ concentration	ppb	—	Uniform(275000, 281000)
CO ₂ forcing in base year (1849)	Wm ⁻²	—	Uniform(0.13, 0.16)
Initial local surface temperature anomaly for regions 1-8	°C	—	Normal(0, 1) truncated between -1 and 1
AR(1) standard deviation (temperature)	°C	—	Uniform(0, 0.2)
AR(1) standard deviation (sea level rise)	m	—	Uniform(0, 0.05)
AR(1) autocorrelation (temperature)	—	—	Uniform(-0.99, 0.99)
AR(1) autocorrelation (sea level rise)	—	—	Uniform(-0.99, 0.99)
CAR(1) white noise process variance (CO ₂)	—	—	Uniform(0, 200)
CAR(1) measure of correlation memory (CO ₂)	—	—	Uniform(0.01, 11.5)

* We sample $\frac{1.2}{1-N(0.62,0.034)}$ for the equilibrium climate sensitivity distribution used for official U.S. SCC estimates.

Descriptions of the uncertain climate model and statistical process parameters. The third column (“Default Model Distributions”) lists PAGE-Climates’s original parameter distributions that would typically be sampled when estimating the SCC. In this column, the equilibrium climate sensitivity distribution is not specific to PAGE-Climate but corresponds to the standard distribution all models sample for official U.S. SCC estimates. The fourth column (“MCMC Prior Distributions”) shows the prior parameter distributions for the Bayesian model calibrations using MCMC.

Table C.4 | Uncertain parameters for SNEASY+BRICK and their prior distributions.

Uncertain Model Parameters	Units	MCMC Prior Distributions
Equilibrium climate sensitivity	°C	Cauchy(3,2) truncated between 0 and 10
Ocean vertical heat diffusivity	cm ² s ⁻¹	Log-normal(1.1, 0.3)
Aerosol radiative forcing scaling factor	—	Triangular(0, 3, 1)
Ocean carbon diffusivity	m yr ⁻¹	Uniform(0, 200)
Respiration temperature sensitivity	—	Uniform(1, 5)
Carbon fertilization factor	—	Uniform(0, 1)
Global ocean-averaged thermal expansion coefficient (α)	kg m ⁻³ °C ⁻¹	Uniform(0.05, 0.3)
Sensitivity of Greenland ice sheet equilibrium volume to changes in temperature	mSLE °C ⁻¹	Uniform(-4.0, -0.001)
Greenland ice sheet equilibrium volume for 0 °C global surface temperature anomaly	mSLE	Uniform(5.888, 8.832)
Temperature sensitivity of Greenland ice sheet exponential decay rate	yr ⁻¹ °C ⁻¹	Uniform(0.0, 0.001)
Exponential decay rate for 0 °C global surface temperature anomaly (Greenland)	yr ⁻¹	Uniform(0.0, 0.001)
Initial glaciers and small ice caps mass-balance sensitivity to global surface temperature	m °C ⁻¹ yr ⁻¹	Uniform(0.0, 0.041)
Exponential parameter for glaciers and small ice caps area to volume scaling	—	Uniform(0.55, 1.0)
Slope of undisturbed Antarctic ice sheet bed before loading	—	KDE of parameter values calibrated to paleoclimatic data
Profile parameter for parabolic ice sheet surface (Antarctic)	m ^{1/2}	KDE of parameter values calibrated to paleoclimatic data
Height of Antarctic ice sheet runoff line above which precipitation accumulates as snow for 0 °C local temperatures	m	KDE of parameter values calibrated to paleoclimatic data
Proportionality constant for the dependency of runoff line height on local Antarctic surface temperature	(m/°C)	KDE of parameter values calibrated to paleoclimatic data
Annual precipitation for for 0 °C local Antarctic temperature	m ice equivalent yr ⁻¹	KDE of parameter values calibrated to paleoclimatic data
Coefficient for exponential dependency of precipitation on Antarctic temperature.	—	KDE of parameter values calibrated to paleoclimatic data
Proportionality constant relating Antarctic ice sheet runoff decrease with height to precipitation.	—	KDE of parameter values calibrated to paleoclimatic data
Proportionality constant for ice flow at the Antarctic ice sheet grounding line	m yr ⁻¹	KDE of parameter values calibrated to paleoclimatic data
Power for the relation of Antarctic ice flow speed to water depth	—	KDE of parameter values calibrated to paleoclimatic data
Undisturbed Antarctic ice sheet bed height at the continent center	m	KDE of parameter values calibrated to paleoclimatic data
Partition parameter for effect of ocean subsurface temperature on Antarctic ice flux	—	KDE of parameter values calibrated to paleoclimatic data
Trigger temperature at which Antarctic ice sheet fast dynamics disintegration begins to occur	°C	KDE of parameter values calibrated to paleoclimatic data
Antarctic ice sheet fast dynamics disintegration rate when temperature is above trigger temperature	m yr ⁻¹	KDE of parameter values calibrated to paleoclimatic data
Antarctic Ocean temperature sensitivity to global surface temperature	°C / °C	KDE of parameter values calibrated to paleoclimatic data
Antarctic ocean temperature for 0 °C global surface temperature anomaly	°C	KDE of parameter values calibrated to paleoclimatic data
Initial global mean surface temperature anomaly	°C	Normal(0, 1)
Initial ocean heat uptake	10 ²² J	Uniform(-100, 0)
Initial atmospheric CO ₂ concentration	ppm	Uniform(275, 281)
Initial atmospheric N ₂ O concentration	ppb	Uniform(264, 282)
Initial sea level rise due to thermal expansion.	m SLE	Uniform(-0.0484, 0.0484)
Initial volume of Greenland ice sheet	m SLE	Uniform(7.16, 7.56)
Initial total volume of glaciers and small ice caps	m SLE	Uniform(0.31, 0.53)
Initial cumulative sea level rise from glaciers and small ice caps	m	Uniform(-0.0536, 0.0791)
Initial sea level rise from Antarctic ice sheet	m	Uniform(-0.04755, 0.05585)
AR(1) standard deviation (temperature)	°C	Uniform(0, 0.2)
AR(1) standard deviation (ocean heat content)	10 ²² J	Uniform(0, 4)
AR(1) standard deviation (glaciers and small ice caps)	m	Uniform(0, 0.0015)
AR(1) standard deviation (Greenland ice sheet)	m	Uniform(0, 0.002)
AR(1) standard deviation (Antarctic ice sheet)	m	Uniform(0, 0.002)
AR(1) standard deviation (global mean sea level)	m	Uniform(0, 0.05)
AR(1) Autocorrelation (temperature)	—	Uniform(-0.99, 0.99)
AR(1) Autocorrelation (ocean heat content)	—	Uniform(-0.99, 0.99)
AR(1) Autocorrelation (glaciers and small ice caps)	—	Uniform(-0.99, 0.99)
AR(1) Autocorrelation (Greenland ice sheet)	—	Uniform(-0.99, 0.99)
AR(1) Autocorrelation (Antarctic ice sheet)	—	Uniform(-0.99, 0.99)
AR(1) Autocorrelation (global mean sea level)	—	Uniform(-0.99, 0.99)
CAR(1) white noise process variance (CO ₂)	—	Uniform(0, 200)
CAR(1) measure of correlation memory (CO ₂)	—	Uniform(0.01, 11.5)

Descriptions of the uncertain climate model and statistical process parameters. The third column (“MCMC Prior Distributions”) shows the prior parameter distributions for the Bayesian model calibrations using MCMC. Uncertain Antarctic ice sheet prior distributions correspond to kernel density estimates of 800,000 parameter samples from a previous calibration to paleoclimate data spanning 240,000 years before the current era to the present.

Table C.5 | Model and scenario specific social cost of carbon estimates.

IAM	RCP8.5					RCP2.6	Remove Parameter Correlations	U.S. Uncertainty	Neglect Sea Level Observations
	Constant 2.5%	Constant 3%	Constant 5%	Ramsey 1.5%, 1.0	Ramsey 1.5%, 1.5				
<i>SB-DICE</i>	96 (58–120)	61 (38–79)	18 (12–23)	63 (37–85)	32 (19–44)	31 (21–42)	64 (33–100)	70 (25–140)	59 (31–90)
<i>SB-FUND</i>	15 (7–25)	10 (4–17)	2 (0–5)	11 (5–19)	7 (2–12)	0 (–4–5)	10 (2–21)	12 (–1–34)	9 (2–16)
<i>SB-PAGE</i>	130 (61–170)	77 (37–100)	16 (9–22)	33 (16–46)	13 (7–18)	17 (10–25)	74 (30–120)	78 (20–200)	71 (28–95)
Column Average	80 (42–105)	49 (26–65)	12 (7–17)	36 (19–50)	17 (9–25)	16 (9–24)	49 (22–80)	53 (15–125)	46 (20–67)

Reported values show the average SCC estimate for a specific scenario-model combination, with lower values in parenthesis showing the corresponding 5-95% predictive credible interval. The five columns under the RCP8.5 heading depict estimates under constant consumption discounting (2.5%, 3%, and 5%) and Ramsey discounting (pure rate of time preference = 1.5%, elasticity of marginal utility of consumption = 1.0 and 1.5). The remaining columns show alternative SCC estimates using a constant 3% consumption discount rate for the RCP2.6 scenario (“RCP2.6”), removing the posterior correlations by independently sampling each marginal posterior distribution (“Remove Parameter Correlations”), neglecting posterior parameter correlations and sampling the equilibrium climate sensitivity distribution specified for official U.S. SCC estimates (“U.S. Uncertainty”), and a scenario where SNEASY+BRICK is recalibrated without sea level observations (“Neglect Sea Level Observations”). The bottom row shows each scenario’s average SCC estimate.

REFERENCES

- Adler, Matthew, David Anthoff, Valentina Bosetti, Greg Garner, Klaus Keller, and Nicolas Treich. 2017. “Priority for the Worse-off and the Social Cost of Carbon.” *Nature Climate Change* 7: 433–49.
- Anthoff, David. 2011. “Optimal Global Dynamic Carbon Abatement.” *Working Paper*, 1–27.
- . 2019a. “FUND Documentation.” 2019.
<https://fundjl.readthedocs.io/en/latest/index.html>.
- . 2019b. *Mimi-DICE-2010.Jl* (version v0.1.2). Julia.
<https://github.com/anthofflab/MimiDICE2010.jl>.
- Anthoff, David, and Johannes Emmerling. 2019. “Inequality and the Social Cost of Carbon.” *Journal of the Association of Environmental and Resource Economists* 6 (2): 243–73.
- Anthoff, David, Frank Errickson, and Lisa Rennels. 2018. *MimiRICE2010.Jl v1.1.0*. Julia.
- Anthoff, David, Cameron Hepburn, and Richard S.J. Tol. 2009. “Equity Weighting and the Marginal Damage Costs of Climate Change.” *Ecological Economics* 68: 836–49.
- Anthoff, David, and Richard S.J. Tol. 2014. “The Climate Framework for Uncertainty, Negotiation and Distribution (FUND), Technical Description, Version 3.9.”
- Arrow, Kenneth J., W. Cline, K.G. Mäler, M. Munasinghe, R. Squitieri, and J. Stiglitz. 1996. “Intertemporal Equity, Discounting, and Economic Efficiency.” Cambridge University Press.
- Arrow, Kenneth J., Maureen L. Cropper, Christian Gollier, Ben Groom, Geoffrey M. Heal, Richard G. Newell, William D. Nordhaus, et al. 2014. “Should Governments Use a Declining Discount Rate in Project Analysis?” *Review of Environmental Economics and Policy* 8 (2): 143–63.
- Bakker, Alexander M. R., Patrick Applegate, and Klaus Keller. 2016. “A Simple, Physically Motivated Model of Sea-Level Contributions from the Greenland Ice Sheet in Response to Temperature.” *Environmental Modeling and Software* 83: 27–35.
- Bakker, Alexander M. R., Tony E. Wong, Kelsey L. Ruckert, and Klaus Keller. 2017. “Sea-Level Projections Representing the Deeply Uncertain Contribution of the West Antarctic Ice Sheet.” *Scientific Reports* 7 (3880): 1–7.
- Bamber, Jonathan L., Michael Oppenheimer, Robert E. Kopp, Willy P. Aspinall, and Roger M. Cooke. 2019. “Ice Sheet Contributions to Future Sea-Level Rise from Structured Expert Judgement.” *Proceedings of the National Academy of Sciences of the United States of America* 116 (23): 11195–11200.

- Bamber, Jonathan L., Riccardo E.M. Riva, Bert L.A. Vermeersen, and Anne M. LeBrocq. 2009. "Reassessment of the Potential Sea-Level Rise from a Collapse of the West Antarctic Ice Sheet." *Science* 324: 901–3.
- Bayes, Thomas, and M. Price. 1763. "An Essay towards Solving a Problem in the Doctrine of Chances. By the Late Rev. Mr. Bayes, F. R. S. Communicated by Mr. Price, in a Letter to John Canton, A. M. F. R. S." *Philosophical Transactions of the Royal Society* 53: 370–418.
- Bollen, Johannes, Bob van der Zwaan, Corjan Brink, and Hans Eerens. 2009. "Local Air Pollution and Global Climate Change: A Combined Cost-Benefit Analysis." *Resource and Energy Economics* 31 (3): 161–81.
- Boucher, O., D. Randall, P. Artaxo, C. Bretherton, G. Feingold, P. Forster, V.M. Kerminen, et al. 2013. "Clouds and Aerosols. In: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change." Cambridge, United Kingdom and New York, NY, USA: Intergovernmental Panel on Climate Change.
- Budolfson, Mark B., David Anthof, Francis Dennig, Frank Errickson, Kevin Kuruc, Dean Spears, and Navroz K. Dubash. 2019. "Minimal Utilitarian Benchmarks for NDCs Yield Equity, Climate, and Development Advantages." *Working Paper*, 1–41.
- Budolfson, Mark B., Francis Dennig, Marc Fleurbaey, Noah Scovronick, Asher Siebert, Dean Spears, and Fabian Wagner. 2018. "Optimal Climate Policy and the Future of World Economic Development." *World Bank Economic Review* 33: 21–40.
- Burke, Marshall, W.M. Davis, and Noah S. Diffenbaugh. 2018. "Large Population Reduction in Economic Damages under UN Mitigation Targets." *Nature* 557: 549–53.
- Burke, Marshall, Solomon M. Hsiang, and Edward Miguel. 2015. "Global Non-Linear Effect of Temperature on Economic Production." *Nature* 527: 235–39.
- Burnett, R., H. Chen, M. Szyszkowicz, N. Fann, B. Hubbell, A. Pope III, J. Apte, et al. 2018. "Global Estimates of Mortality Associated with Long-Term Exposure to Outdoor Fine Particulate Matter." *Proceedings of the National Academy of Sciences of the United States of America* 115: 9592–97.
- California Air Resources Board. 2017. "California's 2017 Climate Change Scoping Plan: The Strategy for Achieving California's 2030 Greenhouse Gas Target."
- Carleton, Tamma A., and Solomon M. Hsiang. 2016. "Social and Economic Impacts of Climate." *Science* 353 (6304).
- Carleton, Tamma A., Amir Jina, Michael T. Delgado, Michael Greenstone, Trevor Houser, Solomon M. Hsiang, Andrew Hultgren, et al. 2020. "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits." *NBER Working Paper* 27599.

- Chang, Kelly M., Jeremy J. Hess, John M. Balbus, Jonathan J. Buonocore, David A. Cleveland, Maggie L. Grabow, Roni Neff, Rebecca K. Saari, Christopher W. Tessum, and Paul Wilkinson. 2017. "Ancillary Health Effects of Climate Mitigation Scenarios as Drivers of Policy Uptake: A Review of Air Quality, Transportation and Diet Cobenefits Modeling Studies." *Environmental Research Letters* 12 (113001).
- Chichilnisky, G., and G. Heal. 1994. "Who Should Abate Carbon Emissions? An International Viewpoint." *Economics Letters* 44 (4): 443–49.
- Church, John A., P.U. Clark, J.M. Cazenave, S. Jevrejeva, A. Levermann, M.A. Merrifield, G.A. Milne, et al. 2013. "Sea Level Change. In: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change." Cambridge, United Kingdom and New York, NY, USA: Intergovernmental Panel on Climate Change.
- Church, John A., and Neil J. White. 2011. "Sea-Level Rise from the Late 19th to the Early 21st Century." *Surveys in Geophysics* 35: 585–602.
- Clarke, L., K. Jiang, K. Akimoto, M. Babiker, G. Blanford, K. Fisher-Vanden, J.C. Hourcade, et al. 2014. "Assessing Transformation Pathways. In: Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change." Cambridge, United Kingdom and New York, NY, USA: Intergovernmental Panel on Climate Change.
- Cohen, Aaron J., Michael Brauer, Richard T. Burnett, H.R. Anderson, Joseph Frostad, Kara Estep, Kalpana Balakrishnan, et al. 2017. "Estimates and 25-Year Trends of the Global Burden of Disease Attributable to Ambient Air Pollution: An Analysis of Data from the Global Burden of Diseases Study 2015." *Lancet* 389: 1907–18.
- Collins, M., R. Knutti, J. Arblaster, J.L. Dufresne, T. Fichefet, P. Friedlingstein, X. Gao, et al. 2013. "Long-Term Climate Change: Projections, Commitments and Irreversibility. In: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change." Cambridge, United Kingdom and New York, NY, USA: Intergovernmental Panel on Climate Change.
- Committee on the Medical Effects of Air Pollutants. 2010. "The Mortality Effects of Long-Term Exposure to Particulate Air Pollution in the United Kingdom: Health Protection Agency Report." London, United Kingdom: Health Protection Agency.
- Dasgupta, Susmita, Benoit Laplante, Siobhan Murray, and David Wheeler. 2011. "Exposure of Developing Countries to Sea-Level Rise and Storm Surges." *Climatic Change* 106: 567–79.
- DeConto, Robert M., and David Pollard. 2016. "Contribution of Antarctica to Past and Future Sea-Level Rise." *Nature* 531: 591–597.

- Dennig, Francis, Mark B. Budolfson, Marc Fleurbaey, Asher Siebert, and Robert H. Socolow. 2015. "Inequality, Climate Impacts on the Future Poor, and Carbon Prices." *Proceedings of the National Academy of Sciences of the United States of America* 112 (52): 15827–32.
- Dennig, Francis, and Johannes Emmerling. 2017. "A Note on Optima with Negishi Weights." *Working Paper*.
- Desaigues, B., D. Ami, A. Bartczak, M. Braun-Kohlova, S. Chilton, M. Czajkowski, V. Farreras, et al. 2011. "Economic Valuation of Air Pollution Mortality: A 9-Country Contingent Valuation Survey of Value of a Life Year (Voly)." *Ecological Indicators* 11 (3): 902–10.
- Di, Qian, Yan Wang, Antonella Zanobetti, Yun Wang, Petros Koutrakis, Christine Choirat, Francesca Dominici, and Joel D. Schwartz. 2017. "Air Pollution and Mortality in the Medicare Population." *The New England Journal of Medicine* 376: 2513–22.
- Diaz, Delavane, and Francis Moore. 2017. "Quantifying the Economic Risks of Climate Change." *Nature Climate Change* 7: 774–82.
- Dietz, Simon, Frederick van der Ploeg, Armon Rezai, and Frank Venmans. 2020. "Are Economists Getting Climate Dynamics Right and Does It Matter?" *University of Oxford Department of Economics Discussion Paper Series*.
- Dlugokencky, E.J. 1994. "The Growth Rate and Distribution of Atmospheric Methane." *Journal of Geophysical Research* 99 (D8): 17021–43.
- Dyrugerov, Mark B., and Mark F. Meier. 2005. "Glaciers and the Changing Earth System: A 2004 Snapshot." Occasional Paper No. 58. Institute of Arctic and Alpine Research, University of Colorado, Boulder.
- Erickson, Frank C., William D. Collins, Vivek Srikrishnan, Klaus Keller, and David Anthoff. 2020. "Improving the Climate Model Calibration Reduces the Social Cost of Methane." *Working Paper*.
- Etheridge, D.M., L.P. Steele, R.J. Francey, and R.L. Langenfelds. 1998. "Atmospheric Methane between 1000 A.D. and Present: Evidence of Anthropogenic Emissions and Climatic Variability." *Journal of Geophysical Research* 103 (D13): 15979–93.
- Etheridge, D.M., L.P. Steele, R.L. Langenfelds, and R.J. Francey. 1996. "Natural and Anthropogenic Changes in Atmospheric CO₂ over the Last 1000 Years from Air in Antarctic Ice and Firn." *Journal of Geophysical Research* 101 (D2): 4115–28.
- Etminan, M., G. Myhre, E. J. Highwood, and K. P. Shine. 2016. "Radiative Forcing of Carbon Dioxide, Methane, and Nitrous Oxide: A Significant Revision of the Methane Radiative Forcing." *Geophysical Research Letters* 43: 12,614–12,623.
- European Commission Joint Research Centre. 2016. "Emission Database for Global Atmospheric Research (EDGAR)." 2016.

- Executive Office of the President. 2017. *Executive Order 13783: Promoting Energy Independence and Economic Growth*. 82 FR 16093.
- Field, C.B., V.R. Barros, D.J. Dokken, K.J. Mach, M.D. Mastrandrea, T.E. Bilir, M. Chatterjee, et al. 2014. “Summary for Policymakers. In: Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part A: Global and Sectoral Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change.” Cambridge, United Kingdom and New York, NY, USA: Intergovernmental Panel on Climate Change.
- Flato, G., J. Marotzke, B. Abiodun, P. Braconnot, S.C. Chou, W. Collins, P. Cox, et al. 2013. “Evaluation of Climate Models. In: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change.” Cambridge, United Kingdom and New York, NY, USA: Intergovernmental Panel on Climate Change.
- Fleurbaey, Marc, Maddalena Ferranna, Mark Budolfson, Francis Dennig, Kian Mintz-Woo, Robert Socolow, Dean Spears, and Stephane Zuber. 2019. “The Social Cost of Carbon: Valuing Inequality, Risk, and Population for Climate Policy.” *The Monist* 102: 84–109.
- Forest, Chris E., Peter H. Stone, Andrei P. Sokolov, Myles R. Allen, and Mort D. Webster. 2002. “Quantifying Uncertainties in Climate System Properties with the Use of Recent Climate Observations.” *Science* 295: 113–17.
- Forestiere, F., H. Kan, and A. Cohen. 2014. “Background Paper 4: Updated Exposure-Response Functions Available for Estimating Mortality Impacts.” Copenhagen, Denmark: World Health Organization Europe.
- Frederikse, Thomas, Maya K. Buchanan, Erwin Lambert, Robert E. Kopp, Michael Oppenheimer, D.J. Rasmussen, and Roderik S.W. van de Wal. 2020. “Antarctic Ice Sheet and Emission Scenario Controls on 21st-Century Extreme Sea-Level Changes.” *Nature Communications* 11 (390).
- Fullerton, Don, and Daniel H. Karney. 2018. “Multiple Pollutants, Co-Benefits, and Suboptimal Environmental Policies.” *Journal of Environmental Economics and Management* 87: 52–71.
- Fuzzi, S., U. Baltensperger, K. Carslaw, S. Decasari, H. Denier van der Gon, M.C. Facchini, D. Fowler, et al. 2015. “Particulate Matter, Air Quality and Climate: Lessons Learned and Future Needs.” *Atmospheric Chemistry and Physics* 15: 8217–99.
- Gelman, Andrew, John B. Carlin, Hal S. Stern, and Donald B. Rubin. 1995. *Bayesian Data Analysis*. United States: Chapman & Hall/CRC.
- Gilks, W.R., S. Richardson, and D.J. Spiegelhalter. 1996. *Markov Chain Monte Carlo in Practice*. London, UK: Chapman & Hall/CRC.

- Gillingham, Kenneth, William D. Nordhaus, David Anthoff, Geoffrey Blanford, Valentina Bosetti, Peter Christensen, Haewon McJeon, and John Reilly. 2018. "Modeling Uncertainty in Integrated Assessment of Climate Change: A Multi-Model Comparison." *Journal of the Association of Environmental and Resource Economists* 5 (4): 791–826.
- Gouretski, Viktor, and Klaus P. Koltermann. 2007. "How Much Is the Ocean Really Warming?" *Geophysical Research Letters* 34 (L01610).
- Green, Donna, Lisa Alexander, Kathy McInnes, John Church, Neville Nicholls, and Neil White. 2010. "An Assessment of Climate Change Impacts and Adaptation for the Torres Strait Islands, Australia." *Climatic Change* 102: 405–33.
- Hänsel, Martin C., Moritz A. Drupp, Daniel J.A. Johansson, Frikk Nesje, Christian Azar, Mark C. Freeman, Ben Groom, and Thomas Sterner. 2020. "Climate Economics Support for the UN Climate Targets." *Nature Climate Change* 10: 781–89.
- Harmesen, Mathijs J. H. M., Detlef P. van Vuuren, Maarten van den Berg, Andries F. Hof, Chris Hope, Volker Krey, Jean-Francois Lamarque, Adriana Marcucci, Drew T. Shindell, and Michiel Schaeffer. 2015. "How Well Do Integrated Assessment Models Represent Non-CO₂ Radiative Forcing?" *Climatic Change* 133: 565–82.
- Hartin, C. A., P. Patel, A. Schwarber, R. P. Link, and B. P. Bond-Lamberty. 2015. "A Simple Object-Oriented and Open-Source Model for Scientific and Policy Analyses of the Global Climate System – Hector v1.0." *Geoscientific Model Development* 8: 939–55.
- Hastings, W.K. 1970. "Monte Carlo Sampling Methods Using Markov Chains and Their Applications." *Biometrika* 57 (1): 97–109.
- Hoeting, Jennifer A., David Madigan, Adrian E. Raftery, and Chris T. Volinsky. 1999. "Bayesian Model Averaging: A Tutorial." *Statistical Science* 14 (4): 382–417.
- Hof, Andries F., Chris Hope, Jason Lowe, Michael D. Mastrandrea, Malte Meinshausen, and Detlef P. van Vuuren. 2012. "The Benefits of Climate Change Mitigation in Integrated Assessment Models: The Role of the Carbon Cycle and Climate Component." *Climatic Change* 113: 897–917.
- Hope, Chris. 2005. "The Climate Change Benefits of Reducing Methane Emissions." *Climatic Change* 68: 21–39.
- . 2006. "The Marginal Impacts of CO₂, CH₄, and SF₆ Emissions." *Climate Policy* 6: 537–44.
- . 2011. "The PAGE09 Integrated Assessment Model: A Technical Description." Cambridge Judge Business School Working Papers. University of Cambridge Judge Business School.

- Hsiang, Solomon M., Robert E. Kopp, Amir Jina, James Rising, Michael Delgado, Shashank Mohan, D.J. Rasmussen, et al. 2017. “Estimating Economic Damage from Climate Change in the United States.” *Science* 356 (6345): 1362–69.
- Ikefuji, Masako, Jan R. Magnus, and Hiroaki Sakamoto. 2014. “The Effect of Health Benefits on Climate Change Mitigation Policies.” *Climatic Change* 126: 229–43.
- Interagency Working Group on Social Cost of Greenhouse Gases. 2016. “Addendum to Technical Support Document on Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866: Application of the Methodology to Estimate the Social Cost of Methane and the Social Cost of Nitrous Oxide.” United States: United States Government.
- Interagency Working Group on the Social Cost of Carbon. 2010. “Technical Support Document: Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866.” United States: United States Government.
- . 2016. “Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866.” United States: United States Government.
- IPCC. 2007. *Changes in Atmospheric Constituents and Radiative Forcing. In Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge, United Kingdom: Cambridge University Press.
- Joos, F., R. Roth, J.S. Fuglestedt, G.P. Peters, I.G. Enting, W. von Bloh, V. Brovkin, et al. 2013. “Carbon Dioxide and Climate Impulse Response Functions for the Computation of Greenhouse Gas Metrics: A Multi-Model Analysis.” *Atmospheric Chemistry and Physics* 13: 2793–2825.
- Keeling, C.D., R.B. Bacastow, A.E. Bainbridge, C.A. Ekdahl, P.R. Guenther, and L.S. Waterman. 1976. “Atmospheric Carbon Dioxide Variations at Mauna Loa Observatory, Hawaii.” *Tellus* 28: 538–51.
- Kirezci, Ebru, Ian R. Young, Roshanka Ranasinghe, Sanna Muis, Robert J. Nicholls, Daniel Lincke, and Jochen Hinkel. 2020. “Projections of Global-Scale Extreme Sea Levels and Resulting Episodic Coastal Flooding over the 21st Century.” *Scientific Reports* 10 (11629): 1–12.
- Knutti, Reto, and Gabriele C. Hegerl. 2008. “The Equilibrium Sensitivity of the Earth’s Temperature to Radiation Changes.” *Nature Geoscience* 1: 735–43.
- Koen, Chris. 2005. “The Analysis of Irregularly Observed Stochastic Astronomical Time-Series - I. Basics of Linear Stochastic Differential Equations.” *Monthly Notices of the Royal Astronomical Society* 361 (3): 887–96.
- Kolstad, C., K. Urama, J. Broome, A. Bruvoll, M. Cariño Olvera, D. Fullerton, C. Gollier, et al. 2014. “Social, Economic and Ethical Concepts and Methods. In: Climate Change 2014:

- Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change.” Cambridge, United Kingdom and New York, NY, USA: Intergovernmental Panel on Climate Change.
- Kriegler, Elmar. 2005. “Imprecise Probability Analysis for Integrated Assessment of Climate Change. PhD Thesis.” Dissertation, Potsdam, Germany: University of Potsdam.
- Lamontagne, J.R., P.M. Reed, G. Marangoni, K. Keller, and G.G. Garner. 2019. “Robust Abatement Pathways to Tolerable Climate Futures Require Immediate Global Action.” *Nature Climate Change* 9: 290–94.
- Le Quéré, C., R. M. Andrew, P. Friedlingstein, S. Sitch, J. Hauck, J. Pongratz, P. A. Pickers, et al. 2018. “Global Carbon Budget 2018.” *Earth System Science Data* 10 (4): 2141–94.
- Lempert, Robert J., David G. Groves, Steven W. Popper, and Steve C. Banks. 2006. “A General, Analytic Method for Generating Robust Strategies and Narrative Scenarios.” *Management Science* 52 (4): 514–28.
- Lenton, Timothy M., Johan Rockström, Owen Gaffney, Stefan Rahmstorf, Katherine Richardson, Will Steffen, and Hans Joachim Schellnhuber. 2019. “Climate Tipping Points - Too Risky to Bet Against.” *Nature* 575: 592–95.
- Li, Mingwei, Da Zhang, Chiao-Ting Li, Kathleen M. Mulvaney, Noelle E. Selin, and Valerie J. Karplus. 2018. “Air Quality Co-Benefits of Carbon Pricing in China.” *Nature Climate Change* 8: 398–403.
- Lieberman-Cribbin, Wil, Christina Gillezeau, Rebecca M. Schwartz, and Emanuela Taioli. 2020. “Unequal Social Vulnerability to Hurricane Sandy Flood Exposure.” *Journal of Exposure Science & Environmental Epidemiology*, 1–6.
- Lynch, Cary, Corinne Hartin, Ben Bond-Lamberty, and Ben Kravitz. 2017. “An Open-Access CMIP5 Pattern Library for Temperature and Precipitation: Description and Methodology.” *Earth System Science Data* 9: 281–92.
- Markandya, Anil, Jon Sampedro, Steven J. Smith, Rita Van Dingenen, Cristina Pizarro-Irizar, Inaki Arto, and Gonzalez-Eguino, Mikel. 2018. “Health Co-Benefits from Air Pollution and Mitigation Costs of the Paris Agreement: A Modelling Study.” *Lancet Planetary Health* 2: e126–33.
- Marten, Alex L. 2011. “Transient Temperature Response Modeling in IAMs: The Effects of over Simplification on the SCC.” *Economics: The Open-Access, Open-Assessment E-Journal* 5.
- Marten, Alex L., Elizabeth A. Kopits, Charles W. Griffiths, Stephen C. Newbold, and Ann Wolverson. 2015. “Incremental CH₄ and N₂O Mitigation Benefits Consistent with the Us Government’s SC-CO₂ Estimates.” *Climate Policy* 15 (2): 272–98.

- Marten, Alex L., and Stephen C. Newbold. 2012. “Estimating the Social Cost of Non-CO₂ Ghg Emissions: Methane and Nitrous Oxide.” *Energy Policy* 51: 957–72.
- Matthews, H. Damon, Nathan P. Gillett, Peter A. Stott, and Kirsten Zickfeld. 2009. “The Proportionality of Global Warming to Cumulative Carbon Emissions.” *Nature* 459.
- Matthey, Astrid, and Björn Bünge. 2019. “Methodenkonvention 3.0 zur Ermittlung von Umweltkosten - Kostensätze.” Dessau-Roßlau, Germany: Umweltbundesamt.
- McNeil, Ben I., Richard J. Matear, Robert M. Key, John L. Bullister, and Jorge L. Sarmiento. 2003. “Anthropogenic CO₂ Uptake by the Ocean Based on the Global Chlorofluorocarbon Data Set.” *Science* 299: 233–39.
- Meehl, Gerald A., Catherine A. Senior, Veronika Eyring, Gregory Flato, Jean-Francois Lamarque, Ronald J. Stouffer, Karl E. Taylor, and Manuel Schlund. 2020. “Context for Interpreting Equilibrium Climate Sensitivity and Transient Climate Response from the CMIP6 Earth System Models.” *Science Advances* 6: 1–10.
- Meinshausen, Malte. 2010. “RCP Concentration Calculations and Data: Final Version, Background Data, Acknowledgements and Further Info.” 2010. <http://www.pik-potsdam.de/~mmalte/rcps/index.htm>.
- Meinshausen, Malte, S. C. B. Raper, and T. M. L. Wigley. 2011. “Emulating Coupled Atmosphere-Ocean and Carbon Cycle Models with a Simpler Model, MAGICC6 – Part 1: Model Description and Calibration.” *Atmospheric Chemistry and Physics* 11: 1417–56.
- Meinshausen, Malte, S.J. Smith, K. Calvin, J.S. Daniel, M.L.T. Kainuma, J.F. Lamarque, K. Matsumoto, et al. 2011a. “The RCP Greenhouse Gas Concentrations and Their Extensions from 1765 to 2300.” *Climatic Change* 109: 213–41.
- . 2011b. “The RCP Greenhouse Gas Concentrations and Their Extensions from 1765 to 2300.” *Climatic Change* 109: 213–41.
- Meng, Xiao-Li, and Wing Hung Wong. 1996. “Simulating Ratios of Normalizing Constants via a Simple Identity: A Theoretical Exploration.” *Statistica Sinica* 6: 831–60.
- Metropolis, Nicholas, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. 1953. “Equation of State Calculations by Fast Computing Machines.” *The Journal of Chemical Physics* 21 (6): 1087–92.
- Millar, Richard J., Zebedee R. Nicholls, Pierre Friedlingstein, and Myles R. Allen. 2017. “A Modified Impulse-Response Representation of the Global near-Surface Air Temperature and Atmospheric Concentration Response to Carbon Dioxide Emissions.” *Atmospheric Chemistry and Physics* 17: 7213–28.
- Mitchell, Timothy D. 2003. “Pattern Scaling - an Examination of the Accuracy of the Technique for Describing Future Climates.” *Climatic Change* 60: 217–42.

- Moore, Frances C., Uris Baldos, Thomas Hertel, and Delavane Diaz. 2017. “New Science of Climate Change Impacts on Agriculture Implies Higher Social Cost of Carbon.” *Nature Communications* 8 (1607).
- Moore, Frances C., and Delavane Diaz. 2015. “Temperature Impacts on Economic Growth Warrant Stringent Mitigation Policy.” *Nature Climate Change* 5: 127–31.
- Moore, Frances C., James Rising, Niklas Lollo, Cecilia Springer, Valeri Vasquez, Alex Dolginow, Chris Hope, and David Anthoff. 2018. “Mimi-PAGE, an Open-Source Implementation of the PAGE09 Integrated Assessment Model.” *Scientific Data* 5 (180187).
- Morice, Colin P., John J. Kennedy, Nick A. Rayner, and Phil D. Jones. 2012. “Quantifying Uncertainties in Global and Regional Temperature Change Using an Ensemble of Observational Estimates: The HadCRUT4 Data Set.” *Journal of Geophysical Research* 117 (D08101).
- Myhre, G., Drew Shindell, F.M. Breon, W. Collins, J. Fuglestedt, J. Huang, D. Koch, et al. 2013. “Anthropogenic and Natural Radiative Forcing. In: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change.” Cambridge, United Kingdom and New York, NY, USA: Intergovernmental Panel on Climate Change.
- Myhre, Gunnar. 1998. “New Estimates of Radiative Forcing Due to Well Mixed Greenhouse Gases.” *Geophysical Research Letters* 25 (14): 2715–18.
- National Academies of Sciences, Engineering, and Medicine. 2017. “Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide.” Washington, D.C.: The National Academies Press.
- Nordhaus, William D. 2010. “Economic Aspects of Global Warming in a Post-Copenhagen Environment.” *Proceedings of the National Academy of Sciences of the United States of America* 107: 11721–26.
- . 2017. “Revisiting the Social Cost of Carbon.” *Proceedings of the National Academy of Sciences of the United States of America* 114: 1518–23.
- Nordhaus, William D., and Joseph Boyer. 2000a. *Warming the World - Economic Models of Global Warming*. Cambridge, Massachusetts: The MIT Press.
- . 2000b. *Warming the World: Economic Models of Global Warming*. Cambridge, United Kingdom: MIT Press.
- Nordhaus, William D., and Paul Sztorc. 2013. “DICE 2013R: Introduction and User’s Manual 2nd Edition.”

- Nordhaus, William D., and Zili Yang. 1996. "A Regional Dynamic General-Equilibrium Model of Alternative Climate-Change Strategies." *The American Economic Review* 86 (4): 741–65.
- Oppenheimer, M., B.C. Glavovic, J. Hinkel, R. van de Wal, A.K. Magnan, A. Abd-Elgawad, R. Cai, et al. 2019. "Sea Level Rise and Implications for Low-Lying Islands, Coasts and Communities. In: IPCC Special Report on the Ocean and Cryosphere in a Changing Climate." Intergovernmental Panel on Climate Change.
- Page, Edward A. 2008. "Distributing the Burdens of Climate Change." *Environmental Politics* 17 (4): 556–75.
- Pruss-Ustun, A., C.C. Mathers, C. Corvalan, and A. Woodward. 2003. "Assessing the Environmental Burden of Disease at National and Local Levels. World Health Organization Report." Geneva, Switzerland: World Health Organization.
- Rao, Narasimha D., Bas J. van Ruijven, Keywan Riahi, and Valentina Bosetti. 2017. "Improving Poverty and Inequality Modelling in Climate Research." *Nature Climate Change* 7: 857–62.
- Rao, Shilpa, Zbigniew Klimont, Joana Leitao, Keywan Riahi, Rita van Dingenen, Lara Aleluia Reis, Katherine Calvin, et al. 2016. "A Multi-Model Assessment of the Co-Benefits of Climate Mitigation for Global Air Quality." *Environmental Research Letters* 11.
- Rao, Shilpa, Zbigniew Klimont, Steven J. Smith, Rita Van Dingenen, Frank Dentener, Lex Bouwman, Keywan Riahi, et al. 2017. "Future Air Pollution in the Shared Socio-Economic Pathways." *Global Environmental Change* 42: 346–58.
- Reilly, J., R. Prinn, J. Harnisch, J. Fitzmaurice, H. Jacoby, D. Kicklighter, J. Melillo, P. Stone, and A. Sokolov. 1999. "Multi-Gas Assessment of the Kyoto Protocol." *Nature* 401: 549–55.
- Ricciuto, Daniel M., Kenneth J. Davis, and Klaus Keller. 2008. "A Bayesian Calibration of a Simple Carbon Cycle Model: The Role of Observations in Estimating and Reducing Uncertainty." *Global Biogeochemical Cycles* 22 (GB2030).
- Ricke, Katherine L., and Ken Caldeira. 2014. "Maximum Warming Occurs about One Decade after a Carbon Dioxide Emissions." *Environmental Research Letters* 9: 1–8.
- Roe, Gerard H., and Marcia B. Baker. 2007. "Why Is Climate Sensitivity so Unpredictable?" *Science* 318: 629–32.
- Ruckert, Kelsey L., Yawen Guan, Alexander M. R. Bakker, Chris E. Forest, and Klaus Keller. 2017. "The Effects of Time-Varying Observation Errors on Semi-Empirical Sea-Level Projections." *Climatic Change* 140: 349–60.

- Saari, Rebecca K., Noelle E. Selin, Sebastian Rausch, and Tammy M. Thompson. 2015. "A Self-Consistent Method to Assess Air Quality Co-Benefits from U.S. Climate Policies." *Journal of the Air and Waste Management Association* 65: 74–89.
- Sasgen, Ingo, Michiel van den Broeke, Jonathan L. Bamber, Eric Rignot, Louise Sandberg Sørensen, Bert Wouters, Zdenek Martinec, Isabella Velicogna, and Sebastian B. Simonsen. 2012. "Timing and Origin of Recent Regional Ice-Mass Loss in Greenland." *Earth and Planetary Science Letters*, 293–303.
- Schelling, Thomas C. 1995. "Intergenerational Discounting." *Energy Policy* 23: 395–401.
- Schwarber, Adria K., Steven J. Smith, Corinne A. Hartin, Benjamin Vega-Westhoff, and Ryan L. Sriver. 2019. "Evaluating Climate Emulation: Fundamental Impulse Testing of Simple Climate Models." *Earth System Dynamics* 10: 729–39.
- Scovronick, Noah, Mark B. Budolfson, Francis Dennig, Marc Fleurbaey, Asher Siebert, Robert H. Socolow, Dean Spears, and Fabian Wagner. 2017. "Impact of Population Growth and Population Ethics on Climate Change Mitigation Policy." *Proceedings of the National Academy of Sciences of the United States of America* 114: 12338–43.
- Scovronick, Noah, Mark Budolfson, Francis Dennig, Frank Errickson, Marc Fleurbaey, Wei Peng, Robert H. Socolow, Dean Spears, and Fabian Wagner. 2019. "The Impact of Human Health Co-Benefits on Evaluations of Global Climate Policy." *Nature Communications*, 1–12.
- Scovronick, Noah, Valeri N. Vasquez, Frank C. Errickson, Francis Dennig, Antonio Gasparrini, Shakoor Hajat, Dean Spears, and Mark B. Budolfson. 2019. "Human Health and the Social Cost of Carbon: A Primer and Call to Action." *Epidemiology* 30 (5): 642–47.
- Sherwood, S., M.J. Webb, J.D. Annan, K.C. Armour, P.M. Forster, J.C. Hargreaves, G. Hegerl, et al. 2020. "An Assessment of Earth's Climate Sensitivity Using Multiple Lines of Evidence." *Reviews of Geophysics*.
- Shiell, Leslie. 2003. "Equity and Efficiency in International Markets for Pollution Permits." *Journal of Environmental Economics and Management* 46: 38–51.
- Shindell, Drew. 2015. "The Social Cost of Atmospheric Release." *Climatic Change* 130: 313–26.
- Shindell, Drew, Greg Faluvegi, Karl Seltzer, and Cary Shindell. 2018. "Quantified, Localized Health Benefits of Accelerated Carbon Dioxide Emissions Reductions." *Nature Climate Change* 8 (291).
- Shindell, Drew, Greg Faluvegi, Michael Walsh, Susan C. Anenberg, Rita Van Dingenen, Nicholas Muller, Jeff Austin, Dorothy Koch, and George Milly. 2011. "Climate, Health, Agricultural and Economic Impacts of Tighter Vehicle-Emission Standards." *Nature Climate Change* 1 (59).

- Shindell, Drew, Johan C. I. Kuylenstierna, Elisabetta Vignati, Rita van Dingenen, Markus Amann, Zbigniew Klimont, Susan C. Anenberg, et al. 2012. “Simultaneously Mitigating Near-Term Climate Change and Improving Human Health and Food Security.” *Science* 335: 183–89.
- Silva, Raquel A., J.J. West, Jean-Francois Lamarque, Drew Shindell, William J. Collins, Stig Dalsoren, Greg Faluvegi, et al. 2016. “The Effect of Future Ambient Air Pollution on Human Premature Mortality to 2100 Using Output from the ACCMIP Model Ensemble.” *Atmospheric Chemistry and Physics* 16: 9847–62.
- Smith, A.F.M., and G.O. Roberts. 1993. “Bayesian Computation via the Gibbs Sampler and Related Markov Chain Monte Carlo Methods.” *Journal of the Royal Statistical Society. Series B (Methodological)* 55 (1): 3–23.
- Smith, Christopher J., Piers M. Forster, Myles Allen, Nicholas Leach, Richard J. Millar, Giovanni A. Passerello, and Leighton A. Regayre. 2018. “FAIR v1.3: A Simple Emissions-Based Impulse Response and Carbon Cycle Model.” *Geoscientific Model Development* 11: 2273–97.
- Smith, Matthew R., and Samuel S. Myers. 2018. “Impact of Anthropogenic CO₂ Emissions on Global Human Nutrition.” *Nature Climate Change* 8: 834–39.
- Smith, S.J., J. van Aardenne, Z. Klimont, R.J. Andres, A. Volke, and S. Delgado Arias. 2011. “Anthropogenic Sulfur Dioxide Emissions: 1850–2005.” *Atmospheric Chemistry and Physics* 11: 1101–16.
- Stanton, E.A. 2011. “Negishi Welfare Weights in Integrated Assessment Models: The Mathematics of Global Inequality.” *Climatic Change* 107: 417–32.
- Stern, N.H. 2007. “The Economics of Climate Change: The Stern Review.” Cambridge, United Kingdom: Cambridge University Press.
- Stohl, A., B. Aamaas, M. Amann, L.H. Baker, N. Bellouin, T.K. Berntsen, O. Boucher, et al. 2015. “Evaluating the Climate and Air Quality Impacts of Short-Lived Pollutants.” *Atmospheric Chemistry and Physics* 15: 10529–66.
- Tanaka, Katsumasa, and Elmar Kriegler. 2007. “Aggregated Carbon Cycle, Atmospheric Chemistry, and Climate Model (ACC2): Description of the Forward and Inverse Modes.” Reports on Earth System Science. Hamburg: Max Planck Institute for Meteorology.
- Tans, Pieter, and Ralph Keeling. 2018. “Trends in Atmospheric Carbon Dioxide: Mauna Loa CO₂ Annual Mean Data.” NOAA Earth System Research Laboratory: Global Monitoring Division. <https://www.esrl.noaa.gov/gmd/ccgg/trends/data.html>.
- The IMBIE Team. 2018. “Mass Balance of the Antarctic Ice Sheet from 1992 to 2017.” *Nature* 558.
- . 2020. “Mass Balance of the Greenland Ice Sheet from 1992 to 2018.” *Nature* 579.

- Thoning, Kirk W., and Pieter P. Tans. 1989. "Atmospheric Carbon Dioxide at Mauna Loa Observatory 2. Analysis of the NOAA GMCC Data, 1974-1985." *Journal of Geophysical Research* 94 (D6): 8549–65.
- Tol, Richard S. J. 2009. "The Economic Effects of Climate Change." *Journal of Economic Perspectives* 23 (2): 29–51.
- Tol, Richard S. J., Roebym J. Heintz, and Petronella E. M. Lammers. 2003. "Methane Emission Reduction: An Application of FUND." *Climatic Change* 57: 71–98.
- Tol, Richard S.J. 1997. "On the Optimal Control of Carbon Dioxide Emissions: An Application of FUND." *Environmental Modeling and Assessment* 2 (3): 151–63.
- United Nations Population Division. 2017. "World Population Prospects: 2017 Revision." United Nations Department of Economic and Social Affairs.
- United States Department of Labor: Bureau of Labor Statistics. 2020. "CPI Inflation Calculator." 2020. https://www.bls.gov/data/inflation_calculator.htm.
- United States Environmental Protection Agency: Office of Air and Radiation & Office of Air Quality Planning and Standards. 2015. "Regulatory Impact Analysis for the Clean Power Plan Final Rule." EPA-452/R-15-003. Research Triangle Park, NJ 27711.
- . 2016. "Regulatory Impact Analysis of the Final Oil and Natural Gas Sector: Emission Standards for New, Reconstructed, and Modified Sources." EPA-452/R-16-002. Research Triangle Park, NJ 27711.
- Urban, Nathan M., Phillip B. Holden, Neil R. Edwards, Ryan L. Sriver, and Klaus Keller. 2014. "Historical and Future Learning about Climate Sensitivity." *Geophysical Research Letters* 41: 2543–52.
- Urban, Nathan M., and Klaus Keller. 2009. "Complementary Observational Constraints on Climate Sensitivity." *Geophysical Research Letters* 36: 1–4.
- . 2010. "Probabilistic Hindcasts and Projections of the Coupled Climate, Carbon Cycle and Atlantic Meridional Overturning Circulation System: A Bayesian Fusion of Century-scale Observations with a Simple Model." *Tellus A* 62: 737–50.
- Van Dingenen, Rita, Frank Dentener, Monica Crippa, Joana Leitao, Elina Marmer, Shilpa Rap, Efisio Solazzo, and Luana Valentini. 2018. "TM5-FASST: A Global Atmospheric Source–Receptor Model for Rapid Impact Analysis of Emission Changes on Air Quality and Short-Lived Climate Pollutants." *Atmospheric Chemistry and Physics* 18: 16173–211.
- Vega-Westhoff, Ben, Ryan L. Sriver, C. A. Hartin, Tony E. Wong, and Klaus Keller. 2020. "The Role of Climate Sensitivity in Upper-Tail Sea Level Rise Projections." *Geophysical Research Letters* 47: 1–8.

- Vega-Westhoff, Ben, Ryan L. Sriver, Corinne A. Hartin, Tony E. Wong, and Klaus Keller. 2019. "Impacts of Observational Constraints Related to Sea Level on Estimates of Climate Sensitivity." *Earth's Future* 7: 1–14.
- Vihola, Matt. 2012. "Robust Adaptive Metropolis Algorithm with Coerced Acceptance Rate." *Statistics and Computing* 22 (5): 997–1008.
- Voldoire, A., D. Saint-Martin, S. Senesi, B. Decharme, A. Alias, M. Chevallier, J. Colin, et al. 2019. "Evaluation of CMIP6 DECK Experiments with CNRM-CM6-1." *Journal of Advances in Modeling Earth Systems* 11: 2177–2213.
- Vousdoukas, Michalis I., Roshanka Ranasinghe, Lorenzo Mentaschi, Theodoros A. Plomaritis, Panagiotis Athanasiou, Arjen Luijendijk, and Luc Feyen. 2020. "Sandy Coastlines under Threat of Erosion." *Nature Climate Change* 10: 260–63.
- Vrontisi, Z., J. Abrell, F. Neuwahl, B. Saveyn, and F. Wagner. 2016. "Economic Impacts of EU Clean Air Policies Assessed in a CGE Framework." *Environmental Science and Technology* 55: 54–64.
- Wagner, Gernot, and Martin L. Weitzman. 2018. "Potentially Large Equilibrium Climate Sensitivity Tail Uncertainty." *Economics Letters* 168: 144–46.
- Waldhoff, Stephanie, David Anthoff, Steven Rose, and Richard S.J. Tol. 2014. "The Marginal Damage Costs of Different Greenhouse Gases: An Application of FUND." *Economics: The Open-Access, Open-Assessment E-Journal* 8: 1–33.
<http://dx.doi.org/10.5018/economics-ejournal.ja.2014-31>.
- West, Jason J., Steven J. Smith, Raquel A. Silva, Vaishali Naik, Yuqiang Zhang, Zachariah Adelman, Meridith M. Fry, Susan Anenberg, Larry W. Horowitz, and Jean-Francois Lamarque. 2013. "Co-Benefits of Mitigating Global Greenhouse Gas Emissions for Future Air Quality and Human Health." *Nature Climate Change* 3: 885–89.
- Weyant, John P., Francisco C. de la Chesnaye, and Geoff J. Blanford. 2006. "Overview of EMF-21: Multigas Mitigation and Climate Policy." *The Energy Journal*, no. Multi-Greenhouse Gas Mitigation and Climate Policy Special Issue: 1–32.
- Wigley, T. M. L., and S. C. B. Raper. 2005. "Extended Scenarios for Glacier Melt Due to Anthropogenic Forcing." *Geophysical Research Letters* 32 (5).
- Wong, Tony E., Alexander M. R. Bakker, and Klaus Keller. 2017. "Impacts of Antarctic Fast Dynamics on Sea-Level Projections and Coastal Flood Defense." *Climatic Change* 144: 347–64.
- Wong, Tony E., Alexander M. R. Bakker, Kelsey Ruckert, Patrick Applegate, Aimee B.A. Slangen, and Klaus Keller. 2017. "BRICK v0.2, a Simple, Accessible, and Transparent Model Framework for Climate and Regional Sea-Level Projections." *Geoscientific Model Development* 10: 2741–60.

- Wong, Tony E., Alexandra Klufas, Vivek Srikrishnan, and Klaus Keller. 2018. “Neglecting Model Structural Uncertainty Underestimates Upper Tails of Flood Hazard.” *Environmental Research Letters* 13.
- World Health Organization. 2018. “Ambient (Outdoor) Air Pollution: Fact Sheet.” 2018. [https://www.who.int/en/news-room/fact-sheets/detail/ambient-\(outdoor\)-air-quality-and-health](https://www.who.int/en/news-room/fact-sheets/detail/ambient-(outdoor)-air-quality-and-health).
- World Health Organization Regional Office for Europe. 2013. “Review of Evidence on Health Aspects of Air Pollution: Revihaap Project. Technical Report.”
- Yin, Peng, Michael Brauer, Aaron Cohen, Richard T. Burnett, Jiangmei Liu, Yunning Liu, Ruiming Liang, et al. 2017. “Long-Term Fine Particulate Matter Exposure and Nonaccidental and Cause-Specific Mortality in a Large National Cohort of Chinese Men.” *Environmental Health Perspectives* 125 (117002).
- Zellner, Arnold, and George C. Tiao. 1964. “Bayesian Analysis of the Regression Model with Autocorrelated Errors.” *Journal of the American Statistical Association* 59 (307): 763–78.
- Zivin, Joshua Graff, and Matthew Neidell. 2014. “Temperature and the Allocation of Time: Implications for Climate Change.” *Journal of Labor Economics* 32 (1): 1–26.
- Zscheischler, Jakob, Seth Westra, Bart J.J.M. van den Hurk, Sonia I. Seneviratne, Phillip J. Ward, Andy Pitman, Amir AghaKouchak, et al. 2018. “Future Climate Risk from Compound Events.” *Nature Climate Change* 8: 469–77.