Virtual Seminar on Climate Economics

Organizing Committee:
Glenn Rudebusch (Brookings Institution)
Michael Bauer (University of Hamburg)
Stephie Fried (Arizona State University)
Òscar Jordà (UC Davis, Federal Reserve Bank of San Francisco)
Fernanda Nechio (Federal Reserve Bank of San Francisco)
Toan Phan (Federal Reserve Bank of Richmond)

Minimax-Regret Climate Policy with Deep Uncertainty in Climate Modeling and Intergenerational Discounting

Stephen J. DeCanio, Charles F. Manski, and Alan H. Sanstad

Ecological Economics, Vol. 201, 2022, https://doi.org/10.1016/j.ecolecon.2022.107552

Introduction

Integrated assessment (IA) models enable quantitative evaluation of the benefits and costs of alternative climate policies.

Policy comparisons are performed by considering a planner who seeks to make optimal trade-offs between the costs of carbon abatement and the damages from climate change.

The planning problem has been formalized as a control problem with these components:

(1) equations coupling GHG emissions and abatement to the accumulation of GHGs in the atmosphere and resulting temperature increases.

(2) a damage function that quantifies economic effects of climate change in terms of the loss of global economic output as a function of temperature increases.

(3) an abatement cost function that expresses the cost of actions to reduce GHG emissions relative to a stipulated baseline emissions trajectory.

Costs and damages are expressed as percentage reductions in gross world product.

The problem is to minimize the costs of abatement and damages over a time horizon.

Studying climate policy as a control problem presumes that a planner knows enough to make optimization feasible, but physical and economic uncertainties abound.

Physical scientists have performed multi-model ensemble (MME) analysis. Lacking a consensus climate model, they have developed multiple models. To cope with inter-model *structural uncertainty*, they compute simple or weighted averages of the outputs of MMEs. Choosing appropriate weights has been problematic.

Economists have estimated multiple damage functions and abatement cost functions. In general, economists have not performed MME analyses that combine multiple functions by averaging. They have reported disparate findings from separate studies.

Manski, Sanstad, and DeCanio (*PNAS*, 2021) framed structural uncertainty in climate modeling as a problem of *partial identification*, generating *deep uncertainty*.

This problem refers to situations in which the underlying mechanisms, dynamics, or laws governing a system are not completely known and cannot be credibly modeled definitively even in the absence of data limitations in a statistical sense.

We proposed use of the *minimax-regret* (MMR) decision criterion to account for deep climate uncertainty in integrated assessment without weighting climate model forecasts.

We developed a theoretical framework for cost-benefit analysis of climate policy based on MMR and we applied it computationally with a simple illustrative IA model.

It is important to recognize deep uncertainty in both the physical and economic components of IA models.

Perhaps the most contentious economic issue has been how a planner should assess the costs and benefits of policies across generations.

In our new paper, we study choice of climate policy that minimizes maximum regret with deep uncertainty regarding both the correct climate model and the appropriate intergenerational assessment of policy consequences.

Economists have long framed intergenerational policy assessment using a discount rate.

They have evaluated climate policies by the present discounted value of the sum of abatement costs and the corresponding damages.

There has been considerable debate about what discount rate to use. The choice is consequential.

Low rates favor policies that reduce GHG emissions aggressively and rapidly.
High rates favor policies that act modestly and slowly.

To express deep uncertainty, we suppose that the appropriate discount rate lies within an interval that covers the spectrum of rates used in the literature.

Our mathematical analysis is a straightforward generalization of M-S-D.

There we supposed that the correct climate model is one of six prominent models in the literature on climate science, whereas the correct economic model is known.

We supposed that a planner compares six policies, each of which chooses an emissions abatement path that is optimal under one and only one of the six climate models.

Regret is the loss in welfare if the model used in policy making is not correct and, consequently, the chosen abatement path is actually sub-optimal.

The MMR rule chooses a policy that minimizes the maximum regret, or largest degree of sub-optimality, across all six climate models.

Here we suppose that the correct climate model is one of the six examined in M-S-D.

We characterize uncertainty about the discount rate by supposing that it takes one of the seven values $\{0.01, 0.02, \ldots, 0.07\}$, a range that covers the rates commonly used.

This range reflects both empirical uncertainty about the future of the economy and normative uncertainty (or perhaps disagreement) about how the current population values the welfare of future generations.

Given joint uncertainty about the climate model and the discount rate, we suppose that a planner compares forty-three policies.

Forty-two policies entail choosing an emissions abatement path that is optimal under one of the {discount rate, climate model} pairs. The remaining one is a passive policy in which the planner chooses no abatement.

The MMR criterion chooses a policy that minimizes maximum regret across all forty-three potential policies.

The MMR analysis points to use of a discount rate of 0.02 for climate policy.

The MMR decision rule keeps the maximum future temperature increase below 2°C for most of the parameter values used to weight costs and damages.

Prevalent Approaches to Climate and Discount-Rate Uncertainty

Averaging Outputs of MMEs of Climate Models

All climate models are based on a specific set of deterministic nonlinear partial differential equations describing large-scale atmospheric dynamics.

Implementation of the equations is subject to numerous practical choices involving discretization, solution methods, and other details.

Some components of the system – such as cloud formation and heat transfer between land surfaces and the atmosphere – are not yet fully understood and must be approximated.

For these reasons, multiple climate models have been developed and are in use, each reflecting different but credible choices in model design and implementation.

Existing models yield different projections of the global climate.

The range of projections produced by different models is a gauge of deep uncertainty about the climate system given the current state-of-the-science.

Virtually all methods of MME analysis combine model outputs into single projections of future climate variables.

However, climate researchers have recognized persistent methodological problems in combining model projections.

A common technique is to take the simple average across model projections of policy-relevant variables.

Researchers may compute weighted average projections when they believe that models can be ranked with respect to relative accuracy.

However, model performance with respect to historical data does not imply skill in predicting the future climate.

Combining MME outputs into single projected trajectories of the future global climate remains a challenging and unresolved problem.

The recent IPCC physical sciences report states:

> "…despite some progress, no universal, robust method for weighting a multi-model projection ensemble is available…"

Uncertainties and Disagreements Regarding the Discount Rate

The economic losses from climate change are represented by damage functions that give the decreases in world-wide output resulting from increases in mean global temperature.

Economists study dynamic optimization by a planner, which entails discounting to quantify the present value of future economic costs and benefits.

The appropriate magnitude of the discount rate has been contentious.

Controversy persists in part because choice of a discount rate is not only an empirical question regarding the future of the economy.

It is also a normative question, concerning social preferences for equity across future generations.

A simple version of the famous *Ramsey formula* provides a transparent expression of the interplay of normative and empirical considerations in choosing a discount rate.

Let the social welfare function be additively separable in the utility of future generations.

Let ρ be the rate at which the planner discounts the utility of future generations.

Let the utility of a representative consumer be an increasing and concave function of consumption, with constant elasticity $(-\eta)$ of marginal utility.

Let g be the annualized growth rate of consumption between time 0 and a future time t.

Ramsey showed that it is optimal to discount future consumption between the present (time 0) and time t at the rate $\delta = \rho + \eta g$.

From the perspective of the present, the empirical value of g may be uncertain. This uncertainty is similar conceptually to the uncertainty that climate modelers face as they attempt to project the future trajectory of climate variables.

ρ formalizes how the planner views intergenerational equity, with $\rho = 0$ if the planner gives equal weight to the welfare of all future generations and $\rho > 0$ if the planner weights welfare more heavily in the near future than in the distant future.

η formalizes the desirability of intergenerational consumption equity.

A planner may feel normative uncertainty about what values of ρ and η to use.

Supposing that the planner aims to represent society, a source of this uncertainty may be normative disagreements within the present population.

Such disagreements were evident in a dispute between Nordhaus (2007), who used the value $\rho = 0.03$, and Stern (2006), who used $\rho = 0.001$.

Stern concluded that policy should seek to reduce GHG emissions aggressively and rapidly. Nordhaus favored policies that act more modestly and slowly.

We argue against any attempt to cope with empirical and normative uncertainty by choosing a single discount rate.

Instead, we study formation of climate policy recognizing a set of possibly appropriate discount rates.

Minimax-Regret Policy Evaluation

To begin, we specify the control problem that a planner would solve with no uncertainty.

The Optimal-Control Problem

Let B_t represent baseline GHG emissions at time t, A_t be GHG abatement actions at time t under some climate policy, measured in the same units as emissions, $C(A_t)$ be the cost of these actions, and $E_t^{A_t} = B_t - A_t$ be the resulting net emissions.

We refer to A_t and $E_t^{A_t}$ as "paths" or "trajectories," and we assume that abatement paths are chosen from some space of feasible paths.

Emissions paths are used as inputs to a climate model M.

We focus on the global mean temperatures projected by M as a function of these paths.

Let $T(E_t^{A_t}, M)$ be the global mean temperature at time t determined by the GHG trajectory $E_t^{A_t}$ when it is predicted by the climate model M.

Then a damage function can be written as $D\left(T\left(E_t^{A_t}, M\right)\right)$.

For abatement path A_t and climate model M, denote the associated total cost (abatement plus damages) at time t as

$$\mathbb{C}(A_t, M) \equiv C(A_t) + D\left(T\left(E_t^{A_t}, M\right)\right)$$

A policymaker seeks to minimize the present value of cost over a planning horizon. As usual in the climate economics literature, we assume an infinite horizon.

The control problem given climate model M is to solve

$$\min_{A_t} \int_0^\infty \mathbb{C}(A_t, M) e^{-\delta t} dt$$

where δ is the discount rate.

We suppose that the optimal A_t is chosen with commitment at time zero. That is, it is not updated over time as new climate or cost information is obtained.

Under certain assumptions, this optimization problem has a unique solution.

The Minimax-Regret Decision Rule

Let $\boldsymbol{\Delta} = \{\delta_1,\ldots, \delta_K\}$ be a set of discount rates and $\mathbf{M} = \{M_1,\ldots,M_N\}$ be a model ensemble.

The planner now faces the problem of minimizing cost over the horizon while recognizing joint {discount rate, model} uncertainty.

For rate δ_i and model M_j, let $A^*_{t;\delta_i,M_j}$ be the optimal abatement path defined by

$$A^*_{t;\delta_i,M_j} = \arg \min_{A_t} \int_0^\infty \mathbb{C}(A_t, M_j)e^{-\delta_i t}\, dt$$

Let $\mathbb{C}^*\left(A^*_{t;\delta_i,M_j}, \delta_i, M_j\right)$ be the associated minimum cost:

$$\mathbb{C}^*\left(A^*_{t;\delta_i,M_j}, \delta_i, M_j\right) = \int_0^\infty \mathbb{C}(A^*_t, M_j)e^{-\delta_i t} dt$$

Now consider any feasible abatement trajectory A_t. The *regret* $\mathbb{R}(A_t, \delta_i, M_j)$ associated with A_t, when discount rate δ_i and climate model M_j describe the world, is the difference between the cost of A_t and the cost of the *optimal* policy associated with δ_i and M_j:

$$\mathbb{R}(A_t, \delta_i, M_j) = \int_0^\infty \mathbb{C}(A_t, M_j) e^{-\delta_i t} dt - \mathbb{C}^* \left(A^*_{t;\delta_i M_j}, \delta_i, M_j\right)$$

To apply the MMR rule, the planner considers each feasible abatement path A_t and finds the model and discount rate combination that maximizes regret, solving the problem

$$\max_{\delta_i, M_j} \mathbb{R}(A_t, \delta_i, M_i)$$

$$= \max_{\delta_i, M_j} \left[\int_0^\infty \mathbb{C}(A_t, M_j) e^{-\delta_i t} dt - \mathbb{C}^*(A^*_{t;\delta_i, M_j}, \delta_i, M_j)\right]$$

The MMR solution is to find A_t to solve the problem

$$\min_{A_t} \left[\max_{\delta_i, M_j} \mathbb{R}(A_t, \delta_i, M_j)\right]$$

Rather than use the MMR rule, one might use the minimax rule, which embodies the principle of preparing for the worst case.

MMR analysis uses information in a more nuanced and less conservative way.

If a climate policy maker selects one model and discount rate from an ensemble and chooses an emissions abatement path that is optimal for that {model, rate} pair, regret is the excess cost of that abatement path if a different (model, rate) pair is the correct one.

Thus, regret measures the potential sub-optimality of policies.

Choosing a policy to minimize maximum regret means choosing one to minimize the maximum degree of sub-optimality across the set of policies under consideration.

Use of **Δ** to Express Empirical and Normative Uncertainty

The term “uncertainty” has usually referred to incomplete knowledge of the empirical environment of a decision maker, called the “state of nature” or the “state of the world.”

This notion of uncertainty applies to incomplete knowledge of the future global temperature, abatement costs, and damages under alternative climate policies.

We also consider uncertainty about the discount rate.

Our use of the set **Δ** to express both empirical and normative uncertainty regarding the discount rate departs from the usual decision-theoretic focus on empirical uncertainty.

Normative uncertainty may have an empirical source, namely incomplete knowledge of the population preferences that a utilitarian planner would seek to maximize.

The planner may face the difficult task of representing a population whose members may not be clear about their time preferences or concern with intergenerational inequalities.

Using **Δ** to express normative uncertainty is a more radical departure from the decision-theoretic norm if normative disagreements exist within the present population.

A segment of the population may strongly value intergenerational equity whereas another segment may be less concerned with the fate of future generations.

Then one may think it necessary to abandon the idealization of a utilitarian planner and replace it with conceptualization of policy making as a non-cooperative political game.

We nonetheless find it attractive to study MMR decision making in this setting.

The MMR rule has some appeal as a broadly acceptable mechanism for policy choice.

Recall that the *regret* of a policy in a specified state of nature measures its degree of sub-optimality in that state, and maximum regret measures the maximum degree of sub-optimality across all states.

Suppose that the members of a heterogeneous present population disagree on what {discount rate, model} should be considered the "true" state of nature.

Then use of the MMR rule to choose policy minimizes the maximum degree of sub-optimality that will be experienced across the population.

Computational Model

To show the consequences of adoption of the MMR decision rule, we present a simple IA model that summarizes the essential economic and physical mechanisms.

The standard in the literature has been to report results about a century into the future.

Analyzing the uncertainty associated with discount rates necessitates attention to longer time horizons.

Phenomena in the more distant future that are negligible in economic terms with high discount rates become salient with low rates.

Model Details

As a simple expression of complex climate dynamics, we use Matthews *et al.* (2009).

They showed that the "carbon-climate response" (CCR), the change in global mean temperature over periods of decades or longer, varies approximately linearly with the increase in cumulative carbon emissions over the same period.

Net cumulative emissions is

$$\mathbf{E}_t^{A_t} = \int_0^t E_t^{A_t} dt = \int_0^t (B_t - A_t) dt$$

There is no requirement that $(B_t - A_t)$ be non-negative. A_t exceeding B_t implies adoption of mitigation measures that yield negative net emissions.

The CCR vary across climate models. The CCR parameter m for model M_j is estimated by determining the model's projected temperature response when driven by a carbon emissions path according to

$$T_t = m_j \mathbf{E}_t^{A_t}, \qquad j = 1, \ldots .6$$

where T henceforth indicates the temperature increase over its initial value at time $t = 0$.

We estimate m_j with historic and projected emissions and temperature data from model j.

Our model ensemble **M** comprises six Earth System Models (ESMs). These ESMs were used in the Climate Model Intercomparison Project Phase 5 (CMIP5).

Table 1

Earth system models used to estimate Carbon-Climate Response (CCR) parameters, with estimated CCR values (℃ per teraton carbon)

Model and model number	*CCR*
1. GFDL-ESM-2G - Geophysical Fluid Dynamics Laboratory Earth System Model version 2G	0.00157
2. BCC-CSM-1 - Beijing Climate Center Climate System Model version 1.1	0.00186
3. FIO-ESM - FIO-ESM - First Institute of Oceanography Earth System Model	0.00194
4. Had-GEM2-ES - Hadley Global Environmental Model 2 - Earth System	0.00229
5. IPSL-CM5A-MR - Institut Pierre Simon Laplace Coupled Model 5A - Medium Resolution	0.00236
6. MIROC-ESM - Model for Interdisciplinary Research on Climate - Earth System Model	0.00244

We specify abatement cost and climate damage functions in quadratic form to implement the IA model as an optimal control problem, allowing for plausible non-linearity in these functions as the abatement effort A_t and the global temperature increase T_t at time t:

$$C(A_t) = \frac{1}{2}\alpha A_t^2$$

$$D(T_t) = \frac{1}{2}\beta T_t^2$$

The quadratic form and value of α are derived from Dietz and Venmans (2019).

The quadratic form and the value of β are taken from Nordhaus and Moffat (2017).

A baseline emissions trajectory B_t is derived from the "Representative Concentration Pathway (RCP) 8.5" scenario in its extended version to year 2500.

This envisions a relatively high growth rate of global carbon emissions from fossil fuel use through the 21st century, followed by a peak plateau period of constant emissions until 2150, and then a decline to a very low level by 2250.

A smoothed functional form having the same general shape as the RCP 8.5 was fitted by nonlinear least squares. The fitted equation for B_t is

$$B_t = \left(\theta t + \frac{B_0}{\exp(\theta\varphi)}\right) \exp\left(-\theta(t-\varphi)\right) .$$

The control problem is to minimize, for a given discount rate and model, the present value of abatement costs plus climate damages over an infinite horizon, subject to the dynamic relationship between cumulative emissions and temperature:

$$\min_{A_t} \int_0^\infty \frac{1}{2}(\alpha A_t^2 + \beta T_t^2)\, e^{-\delta t} dt$$

subject to

$$\frac{d}{dt}\mathbf{E}_t^{A_t} = E_t^{A_t} = B_t - A_t$$

$$T_t = m\mathbf{E}_t^{A_t}$$

$$\mathbf{E}_0^{A_t} = \mathbf{E}_0$$

The last equation specifies an initial condition for net cumulative emissions.

First-order necessary conditions include two coupled differential equations in abatement and the atmospheric greenhouse gas concentration associated with the optimal abatement:

$$\frac{dA_t}{dt} = \delta\, A_t - \frac{\beta m^2}{\alpha}\, \mathbf{E}_t^{A_t}$$

$$\frac{d\mathbf{E}_t^{A_t}}{dt} = B_t - A_t$$

These equations can be solved in closed form for A_t and $\mathbf{E}_t^{A_t}$.

The model satisfies convexity properties implying that the first-order conditions are sufficient for these to be unique optimal solutions to the control problem.

The left-hand panel of Figure 1 shows the baseline B_t and the optimal abatement A_t for a particular set of parameters.

The right-hand panel shows net cumulative emissions under A_t and under a policy of no abatement ($A_t = 0$ for all t).

Figure 1 — Trajectories of B_t, optimal A_t, $\mathbf{E}_t^{A_t}$, and $\mathbf{E}_t^{\text{noabatement}}$ for $m = 0.002286$, $\alpha = 0.000125$, $\beta = 0.018$, and $\delta = 0.05$

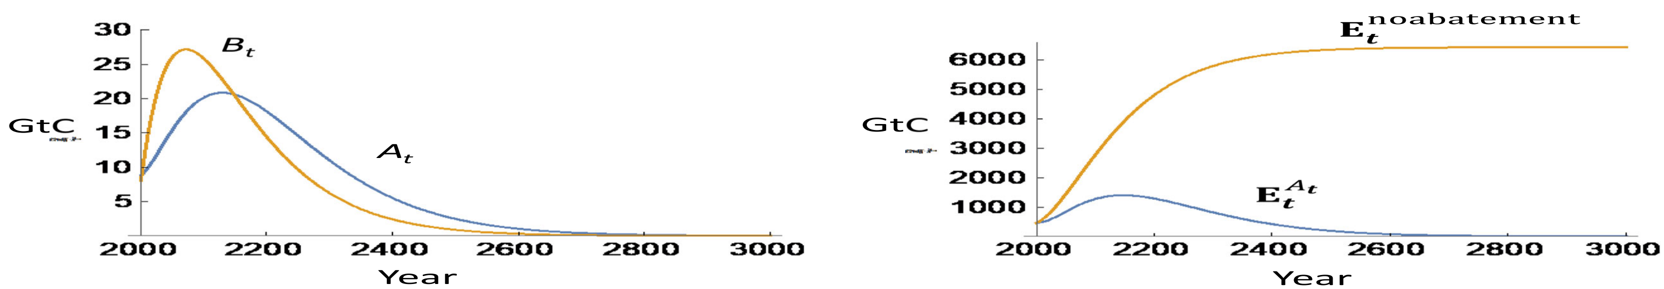

MMR Analysis

We discussed our climate model ensemble **M** above. We specify the possible discount rates as $\mathbf{\Delta} = \{0.01, 0.02, 0.03, 0.04, 0.05, 0.06, 0.07\}$. There are thus forty-two combinations of (δ, m) expressing the range of deep uncertainty.

Regrets can be calculated for any feasible abatement path A_t.

To keep calculation tractable, and because a planner may restrict attention to policies that are optimal in some state of nature, we consider policies that are optimal for some (δ, m) in $\mathbf{\Delta}$ x **M**. There are 42 such policies, and a 43rd when "No Abatement" is a possibility.

To explore the sensitivity of the MMR policy to the α and β parameters, we calculated the MMR for nine combinations of α and β.

Table 2

Values of MMR, uncertain Model and δ, for combinations of α and β
Potential values of δ: {0.01, 0.02, 0.03, 0.04, 0.05, 0.06, 0.07}

Baseline RCP 8.5 (fitted): $B_t = (\theta t + \frac{B_0}{\exp(\theta\varphi)}\exp(-\theta(t-\varphi))\,;\ \theta \to 0.0123125,\ \varphi \to 339.565$

α	β	Model	δ	MMR
0.000075	0.014	IPSL	0.02	0.172
0.000075	0.018	HAD	0.02	0.172
0.000075	0.022	HAD	0.02	0.178
0.000125	0.014	MIROC	0.02	0.266
0.000125	0.018	IPSL	0.02	0.273
0.000125	0.022	IPSL	0.02	0.284
0.0002	0.014	MIROC	0.02	0.478
0.0002	0.018	MIROC	0.02	0.436
0.0002	0.022	MIROC	0.02	0.423

For all (α, β) combinations, the discount rate corresponding to the MMR solution is 0.02.

The IA model allows for calculation of the maximum temperature increase that will be reached for any policy path, and how long it will take to reach that temperature.

Because the actual state of the world is unknown, the temperature increase under the MMR policy cannot be known at the time the policy decision is made.

What is known is that it will be less than or equal to the maximum over all six models, which will occur if MIROC is the true model because m_6 is the greatest of the CCRs.

We find that the MMR decision rule keeps the maximum future temperature increase below 2°C above the 1900-09 level for most parameter values.

Table 3

Values of Maximum Temperature Increase (Tmax) in °C and Years after 2000 when reached,

α = 0.000075 β = 0.014			***α = 0.000075 β = 0.018***			***α = 0.000075 β = 0.022***		
MMR Model	**Years**	**Tmax**	**MMR Model**	**Years**	**Tmax**	**MMR Model**	**Years**	**Tmax**
IPSL	124	1.248	HAD	121	1.056	HAD	118	0.879

α = 0.000125 β = 0.014			***α = 0.000125 β = 0.018***			***α = 0.000125 β = 0.022***		
MMR Model	**Years**	**Tmax**	**MMR Model**	**Years**	**Tmax**	**MMR Model**	**Years**	**Tmax**
MIROC	134	1.831	IPSL	130	1.564	IPSL	125	1.315

α = 0.0002 β = 0.014			***α = 0.0002 β = 0.018***			***α = 0.0002 β = 0.022***		
MMR Model	**Years**	**Tmax**	**MMR Model**	**Years**	**Tmax**	**MMR Model**	**Years**	**Tmax**
MIROC	149	2.660	MIROC	141	2.187	MIROC	135	1.859

Discussion

The MMR rule provides a reasonable way to form climate policy with empirical uncertainty about the climate and normative uncertainty regarding the discount rate.

Our computational analysis offers a new reason for using a low discount rate in climate policy analysis, on the order of 2% per annum.

This discount rate encompasses the pure rate of time preference, intergenerational inequality aversion, projection of the economy's future rate of growth, and other factors that potentially can affect the discount rate.

MMR decision making copes with deep uncertainty without adopting the extreme conservatism of minimax decisions.

MMR enables a planner to deal with heterogeneous populations, who may not themselves be clear about their time preferences or concern with intergenerational equity.

There is no scientific or economic reason that everyone should hold the same normative values. Some people may have only a vague understanding of discounting.

We also find it appealing to view MMR as a consensus-building mechanism.

Calculating regrets enables people with different values to see how implementation of alternative policies might play out from their perspectives.

Our IA model is simple and computationally tractable.

This is partially because we have not considered all possible sources of uncertainty.

The appropriate baseline emissions path is highly uncertain.

The abatement cost and climate damage functions are also uncertain.

We have addressed this partially by sensitivity analysis, calculating MMR solutions with various parameters (α, β) on abatement cost and climate damages.

It would be desirable to expand the analysis to encompass deep uncertainty about the correct values for these weights, a formidable computational task.