



Adaptive mitigation strategies hedge against extreme climate futures

Giacomo Marangoni^{1,2} · Jonathan R. Lamontagne³ · Julianne D. Quinn⁴ · Patrick M. Reed⁵ · Klaus Keller⁶

Received: 10 December 2020 / Accepted: 17 May 2021 / Published online: 7 June 2021
© The Author(s) 2021

Abstract

The United Nations Framework Convention on Climate Change agreed to “strengthen the global response to the threat of climate change, in the context of sustainable development and efforts to eradicate poverty” (UNFCCC 2015). Designing a global mitigation strategy to support this goal poses formidable challenges. For one, there are trade-offs between the economic costs and the environmental benefits of averting climate impacts. Furthermore, the coupled human-Earth systems are subject to deep and dynamic uncertainties. Previous economic analyses typically addressed either the former, introducing multiple objectives, or the latter, making mitigation actions responsive to new information. This paper aims at bridging these two separate strands of literature. We demonstrate how information feedback from observed global temperature changes can jointly improve the economic and environmental performance of mitigation strategies. We focus on strategies that maximize discounted expected utility while also minimizing warming above 2 °C, damage costs, and mitigation costs. Expanding on the Dynamic Integrated Climate-Economy (DICE) model and previous multi-objective efforts, we implement closed-loop control strategies, map the emerging trade-offs and quantify the value of the temperature information feedback under both well-characterized and deep climate uncertainties. Adaptive strategies strongly reduce high regrets, guarding against mitigation overspending for less sensitive climate futures, and excessive warming for more sensitive ones.

Keywords Climate risk management · Adaptive mitigation pathways · Integrated assessment modelling · Multi-objective optimization · Direct policy search

1 Introduction

What is an optimal course of climate mitigation action? Two critical aspects hinder the ability of climate-economy models to provide definitive quantitative answers. First,

✉ Giacomo Marangoni
giacomo.marangoni@polimi.it

deep uncertainties surround key system properties such as climate sensitivity and climate damages (Wagner and Zeckhauser 2018; Knutti et al. 2017; Weitzman 2012; Hwang et al. 2013). Uncertainties are deep in the sense that experts do not agree on a single probability distribution to characterize them (Kwakkel et al. 2016). Assessments neglecting the deeply uncertain nature of these uncertainties and their tails may provide overconfident and misleading results (Pindyck 2013). Second, even when agreeing on the distribution of alternative futures, decision-makers might rank mitigation strategies differently depending on their preferences (Garner et al. 2016). Specific choices of utility, including discounting and risk aversion assumptions, can highly influence the value of costs and benefits of mitigation (Heal 2017; Pindyck 2013; Cai et al. 2016).

Shifting climate-economy models to a broader risk-management approach is a key avenue towards providing a more robust support to climate policy-makers (Heal 2017; Kunreuther et al. 2013; Kwakkel et al. 2016). For one, this approach can help stakeholders to manage the trade-offs of conflicting objectives (UNFCCC 2015; Garner et al. 2016). A multi-objective decision framework offers numerous advantages over single-objective ones: it helps stakeholders better navigate trade-offs (Hobbs et al. 1992), it supports a-posteriori weighing of objectives to explore a wider range of decision makers' preferences, it can account for non-convexities (Goes et al. 2011) and it can lead to the discovery of solutions more robust under deeply uncertain future conditions (Herman et al. 2015).

Previous studies have been limited in their exploration of the synergies and trade-offs between multiple objectives, the future opportunities to revise action, and the strategies to hedge against deep uncertainties. The most commonly discussed mitigation strategies consist of predetermined century-long CO₂ reduction schedules that maximize a single objective (Nordhaus 2017). This single objective is typically maximizing expected global utility, derived from a weighted sum of per-capita consumption discounted over time and averaged across climate futures. Recent analyses have started to examine trade-offs between other objectives (Stechow et al. 2016; Garner et al. 2016). Still, they miss the value of adapting mitigation actions as the climate evolves to improve performance across multiple objectives.

Designing adaptive climate mitigation strategies is an area of active research. One strand of literature analyzes how a risk-averse policy maker copes with the increasing probabilities of extreme climate events as temperature increases (Cai et al. 2016; Cai and Lontzek 2018). Optimal mitigation trajectories stem from recursively evaluating climate action against these changing probabilities. This approach can quickly become computationally intractable (Giuliani et al. 2015). Either uncertainties are resolved at an overly reduced number of predetermined points in time (Urban and Keller 2010), or some approximations are introduced (Shayegh and Thomas 2015; Cai et al. 2016).

An alternative approach adopts policy rules. Policy rules link the control to the state of the system and provide a convenient formulation for suggesting actions that are conditional on observations (Haustein et al. 2017). Previous studies have introduced closed-loop control for climate mitigation within stylized climate models, using single-objective optimization and assuming well-characterized uncertainty (Jarvis et al. 2009; Jarvis et al. 2012; MacMartin et al. 2014).

This paper contributes to the analysis and potential improvement of strategies for controlling global anthropogenic CO₂ emissions. Starting with the multi-objective approach of (Garner et al. 2016), we introduce three main advances to the literature discussed above: (i) we implement closed-loop control, conditioning action to observed warming, in a full-fledged coupled climate-economy model; (ii) we map out the trade-offs between multiple

performance metrics when relying on adaptive strategies; (iii) we quantify the additional robustness of adaptive strategies over non-adaptive ones under consideration of key deep uncertainties.

We follow the Many Objective Robust Decision Making (MORDM) framework (Kasprzyk et al. 2013). First, we identify an approximate set of Pareto-optimal adaptive abatement strategies, optimizing the multi-objective outcome of the Dynamic Integrated model of Climate and the Economy (DICE) coupled with a climate system emulator (Nordhaus 2017; Krieglner 2005). Objectives are evaluated in expected value against a reference ensemble of States-of-the-World (SOWs). These correspond to plausible future climate responses obtained through Bayesian calibration and sampling of key parameters in a climate system emulator (Wong et al. 2017). Next, the resulting Pareto solutions are re-evaluated for two alternative priors for the climate calibration as well as several alternative assumptions for the damage function. Corresponding changes in the objective space are quantified and compared with non-adaptive counterparts.

2 Methods

2.1 The DICE simulator

We model the relationship between economic growth, greenhouse gas (GHG) emissions and temperature impacts with the Dynamic Integrated model of Climate and the Economy, or DICE, for its transparency, computational efficiency and reputation (see Supplementary Information for learning more about DICE and the reasons behind this choice). We depart from the standard welfare-maximizing DICE setup by reformulating the model as a simulator. The simulator evaluates equations iteratively over time, starting in 2015 and stopping in 2250, with 5-year time steps. At each time step DICE receives directions on how many tons of CO₂ to abate from 2020 onward, assuming fixed baseline rates on how much final good to invest for future economic growth.

A previous multi-objective optimization analysis already implemented DICE as a simulator (Garner et al. 2016). Our implementation differs in two ways. First, we consider the parametric assumptions of the last available iteration of the model, DICE-2016R (Nordhaus 2017), instead of the previous DICE-2013R. This explains most of the differences in the two resulting Pareto fronts derived in this and the previous multi-objective study (the remaining changes are due to a different climate model calibration), and results in higher climate risks, as also shown in a more recent sensitivity analysis (Lamontagne et al. 2019). Second we implement the DICE code in the Python programming language, with potential advantages for rapid prototyping.

We focused on the first two centuries as done in recent papers involving a similar simulation approach (Lamontagne et al. 2019; Garner et al. 2016), with the aim of informing global mitigation pathways especially before 2100. While a longer time horizon might change objective values (and discounted ones to a lesser extent), we would still expect our conclusions about the effectiveness of our adaptive approach to hold. In particular, adaptive strategies can be simulated past the horizon they are optimized to without exhibiting end-of-time behavior, since the decision does not depend on time. An open loop strategy would require a longer optimization horizon to know what to do further into the future, with possible undesired end-of-time effects.

2.2 The climate emulator

We replace DICE default climate module with the Diffusion Ocean Energy balance CLIMate model (DOECLIM) (Kriegler 2005) to better capture the non-linearities and uncertainties of the climate system. This choice is particularly relevant for this work, as we aim to improve the assessment of costs and benefits of mitigation, and we focus on a feedback loop involving the climate response itself. DOECLIM is an atmosphere-ocean-land energy balance model. Its core equations derive from a linear approximation of the climate response to perturbations in the Earth's energy budget. The model groups relevant components of the climate into four boxes: land, troposphere over land, troposphere over the sea and ocean mixed layer. Their dynamics are coupled with a one-dimensional upwelling-diffusion ocean model, which represents the heat uptake of the ocean and drives a major part of the transient response.

Uncertainty in the climate response is accounted for by means of three DOECLIM parameters: equilibrium climate sensitivity, vertical ocean diffusivity, and aerosol scaling. Equilibrium climate sensitivity quantifies the increase in atmospheric temperature due to a doubling of CO₂ concentrations. The increase is relative to pre-industrial conditions, and refers to the period after which the climate reaches its new equilibrium. Vertical ocean diffusivity controls the global mean rate of ocean heat uptake. Smaller (larger) diffusivities result in a slower (faster) mixing of heat into the deep ocean. The aerosol scaling factor multiplies the aerosol forcing, reflecting the large uncertainty in the magnitude of this component. We perform a joint Bayesian calibration of the distributions of these parameters within the Building blocks for Relevant Ice and Climate Knowledge (BRICK) framework (Wong et al. 2017), as documented in the Supplementary Information.

2.3 Mitigation strategies

We define mitigation strategies as rules prescribing reductions in global CO₂ emissions from fossil fuels and industry. We focus on abated emissions rather than policies, acknowledging that the same strategy can be implemented with different policy instruments (e.g., carbon tax, or carbon trade) and burden sharing schemes across regions. We distinguish between adaptive and non-adaptive mitigation strategies. In non-adaptive strategies, abatement at each time step is decided a priori, and remains unchanged regardless of the observed conditions. Decision variables are as many as the time steps considered. In adaptive strategies, abatement at each time step and SOW depends on the recently observed atmospheric temperatures. Decision variables are as many as the parameters shaping the function that relates temperatures to abatement.

Three constraints apply to the deployment of abatement. First, we limit the increase in emission reductions to four percentage absolute points per year (Grübler et al. 1999; Keller et al. 2008). Thus, the energy system can become carbon neutral approximately in the year 2040, but not earlier. Second, we assume that abatement expenditures represent long-term decisions in the decarbonization of the energy system. Hence we require abatement to not decrease over time, until the system is carbon neutral. Third, we exclude availability of negative emissions before 2150. This choice is the DICE default assumption (Nordhaus and Sztorc 2013) and also reflects the discussions about excessive reliance on negative emissions (Fuss et al. 2014; Vuuren et al. 2017). After 2150, the maximum amount of negative emissions is 20% of the baseline emissions, and is represented by levels of abatement above 100%. For these levels of abatement we drop the requirement of being monotonic as in

the case of abatement below 100%, allowing for the possibility of using negative emissions temporarily.

For adaptive strategies, we parameterize the feedback between state and control with a convenient functional form (see Supplementary Information for details). The choice of temperature to inform abatement follows the focus on uncertainty in the climate response. While the parameters of the climate are not directly observable, they are partially revealed through the higher or lower temperatures realized over time from a given emissions pathway (Urban et al. 2014). We consider two input variables to determine abatement adaptively in a given time step: temperature level in the previous period, and its change from one period before.

Although we assume the uncertain climate parameters to be fixed, and the system to lack any random shocks or other time-varying stochastic processes, adaptivity of solutions remains valuable. By excluding the possibility of inverting the climate model to pin down its true values, our setup reflects the challenges of learning currently unknown parameters like the climate sensitivity. The latter is still described by the IPCC with the same uncertainty range that was given 40 years ago (Knutti et al. 2017), and will still take time to be learned accurately (Urban et al. 2014). Furthermore, by deterministically optimizing to the fixed parameters we believe to be true, we show that high regrets could follow. Our adaptive approach responds to observations, resulting in different decisions being taken in different possible worlds, without assuming what that world will be.

In general, adaptive strategies should be responsive to the whole system state, which includes capital in the case of DICE. For computational tractability, and as a first proof of concept, we focused on strategies responsive only to temperature and temperature change, in line with the literature that envisions mitigation action as informed by observed/anticipated intensity of anthropogenic climate interference (Haustein et al. 2017; Jarvis et al. 2012). Under the default DICE damage and abatement cost curves, capital is mostly driven by the savings rate, which here is exogenously assigned, and it is not subject to any direct random shock or uncertainty. It might be the case that facing different capital levels with the same climate signal could justify higher or lower abatement efforts, e.g., under large damage functions, or alternative socio-economic assumptions, aspects which deserve further exploration in future research.

2.4 Multiple objectives

We seek Pareto-optimal strategies that (i) maximize utility, (ii) minimize two-degree-years, (iii) minimize damage cost and (iv) minimize mitigation cost. The choice of objectives reflects relevant priorities of characteristic climate stakeholders. We start with considering the standard framework of economic actors that aim to maximize utility. We augment this objective to approximate potential preferences of climate change deniers and myopic political actors that strive to minimize costs related to a cause they do not deem important, namely mitigation. We then add objectives to approximate the concerns of those affected by climate impacts, e.g., inhabitants of low-lying land facing future sea-level rise, who are most responsive to the sheer amount of damages ahead, expressed either in geophysical terms (i.e., overshoot above 2 °C), or in economic terms (i.e., discounting the future and accounting for population growth as with CBGE). Choosing any set of three out of four objectives and assuming their levels are fixed, multiple strategies can exhibit different performance in the fourth dimension, which justifies keeping all objectives in the optimization. Including utility also serves the purpose of contrasting the traditional utility-maximization

with a multi-objective approach, which has been shown to help stakeholders better navigate cost-benefit trade-offs (Hobbs et al. 1992).

Warming above 2 °C quantifies how far off a strategy is from achieving a 2 °C temperature target as discussed in the Paris Agreement (UNFCCC 2015). It is measured in two-degree years, using the integral of the time series of future temperatures (with respect to pre-industrial levels) that lie above 2 °C. We stop the integral in the year 2200, focusing on overshoots occurring within the next two centuries. The integral is then averaged across the sampled SOWs. The other objectives represent equivalent expected losses in today's consumption due to either damage costs (iii), abatement costs (iv), or both (i), expressing them in terms of Certainty and Balanced Growth Equivalent (CBGE) (Garner et al. 2016; Anthoff and Tol 2009) (see the Supplementary for the corresponding equations). For utility, we normalize the corresponding CBGE loss so that zero and one represent, respectively, the minimum and maximum expected utility values attainable within the space explored. Higher values of utility are preferred, in alignment with utility-maximizing models.

2.5 Multi-objective evolutionary optimization

We perform the multi-objective optimization by exploring the feasible space of strategies and excluding those which can be improved in one objective without sacrificing the others. As the search evolves, the front of surviving solutions, i.e., the Pareto front, moves closer to the true set of Pareto-optimal solutions.

A common assumption is to fix a computational budget and seek the best Pareto that can be found within this budget, in our case five million model evaluations for each type of strategy. We track convergence throughout the search using hypervolume, a measure of the objective space dominated by the current best estimate of the Pareto front. We compute this metric across several initialization seeds of the optimization algorithm, to ensure that repeated runs converge to similar hypervolumes and the algorithm is not finding local optima. The leveling off of this metric across seeds and iterations, as shown in Figure S2 in the Supplementary, supports the claim for convergence of the algorithm.

We perform the search with Borg, a state-of-the-art auto-adaptive evolutionary algorithm for multi-objective optimization (Hadka and Reed 2012). This algorithm improves on past designs with measures that avoid stagnation in the search for solutions or degeneration in the diversity of solutions. We also rely on its ability to parallelize model evaluations, speeding up convergence almost linearly with the number of processors used, by relying on its master-worker configuration. Further details on its parameterization and use can be found in Hadka and Reed (2012).

2.6 Robustness to deep uncertainty

Although the optimization is carried out assuming a well-characterized uncertainty for climate sensitivity and a deterministic assumption for the damage function, these are two major sources of deep uncertainty (Knutti et al. 2017; Weitzman 2011).

We consider three alternative non-linear damage functions. One is that used in DICE (referred to as “Low”). The second one is similar to the DICE nominal response for temperatures below 2 °C, but increases more steeply at higher temperatures (Weitzman 2012) and is referred to as “Medium”. A sudden increase in the steepness of the damage response mimics the feature of a possible tipping point (Lenton and Ciscar 2013). Finally, we include a more pessimistic assessment at all temperatures [35], referred to as “High”. In this case, a temperature increase of 4 °C can cause more than 30% losses in GDP (Goes et al. 2011).

This is comparable to a central case of more recent econometric assessments where 4 °C warming is associated with damages somewhere between 10 and 70% of GDP, depending on the empirical specification (Burke et al. 2015). Although we did not include higher damage functions or more severe tipping points, in our multi-objective setting an aversion to such levels of damages is well captured via the two-degree-year metric.

For the climate response, we synthesize two alternative distributions for the three calibrated climate parameters (i.e., equilibrium climate sensitivity, vertical ocean diffusivity and aerosol scaling). These calibration scenarios are representative of a lower and higher assessment of climate sensitivity with respect to the nominal case. The lower assessment, referred to as “Low”, results from a more narrow and optimistic climate sensitivity prior (Chylek and Lohmann 2008). The higher assessment, referred to as “High”, results from a wider prior in climate sensitivity (Urban and Keller 2010). Like before, we subsample the posterior distributions of the three parameters of interest to obtain the final SOW ensembles.

We test the robustness of Pareto-optimal adaptive strategies by studying how these solutions can guard against less likely but potentially disruptive climate impacts. We focus on ten different levels of abatement costs, representing higher or lower budgets allocated to climate mitigation, and pick just as many corresponding abatement strategies, both adaptive and non-adaptive. We re-evaluate these strategies against the alternative deeply uncertain scenarios, given by a calibration ensemble and a damage function. We extract one performance metric of interest: the 95th percentile of damage costs. This choice serves our purpose of testing strategies against extreme climate scenarios. We then focus on the regret of not adapting to observed climate change by taking the difference between the two types of solutions. We then present and discuss the resulting robustness values (McPhail et al. 2018).

Solutions are evaluated against a reference ensemble of climate SOWs, the one most in line with the latest IPCC assessment (Knutti et al. 2017), and a vanilla parameterization of the model. The choice of this reference might bias the attainable robustness after re-evaluation (Bartholomew and Kwakkel 2020; Watson and Kasprzyk 2017). On one hand, exposing the solver to these scenarios during the optimization might lead to better adaptive and non-adaptive strategies. On the other, we avoid the risk of over-tuning our solutions to our imagined futures, and offer a quantification of performance under unexpected conditions, which often occur in reality (Fukuyama 2007). We leave the exploration of how multi-scenario optimization could improve robustness for further research.

3 Results and discussion

The results of our multi-objective optimization reveal strong trade-offs between mitigation and damage costs (two right axes of Fig. 1). A solution that maximizes utility (B in Fig. 1) compromises between these two objectives. However, this solution represents one of many possible weightings of mitigation and damage costs. Decision-makers often have trouble interpreting utility functions and may choose a different option than their supposed utility-maximizing policy when presented with other alternatives (Hobbs et al. 1992). While this utility function finds a solution closer to the minimum mitigation cost solution than the minimum damage cost solution, a decision-maker may prefer to further reduce damages upon seeing these trade-offs.

Minimizing damages coincides with minimizing 2 degree-years. A solution that minimizes two-degree years implements as much abatement as possible in both the near- and long-term (solution A in Fig. 1). The energy system becomes carbon neutral by mid-century, mostly bounded by the inertia in the turnover of human and physical capital that limits the

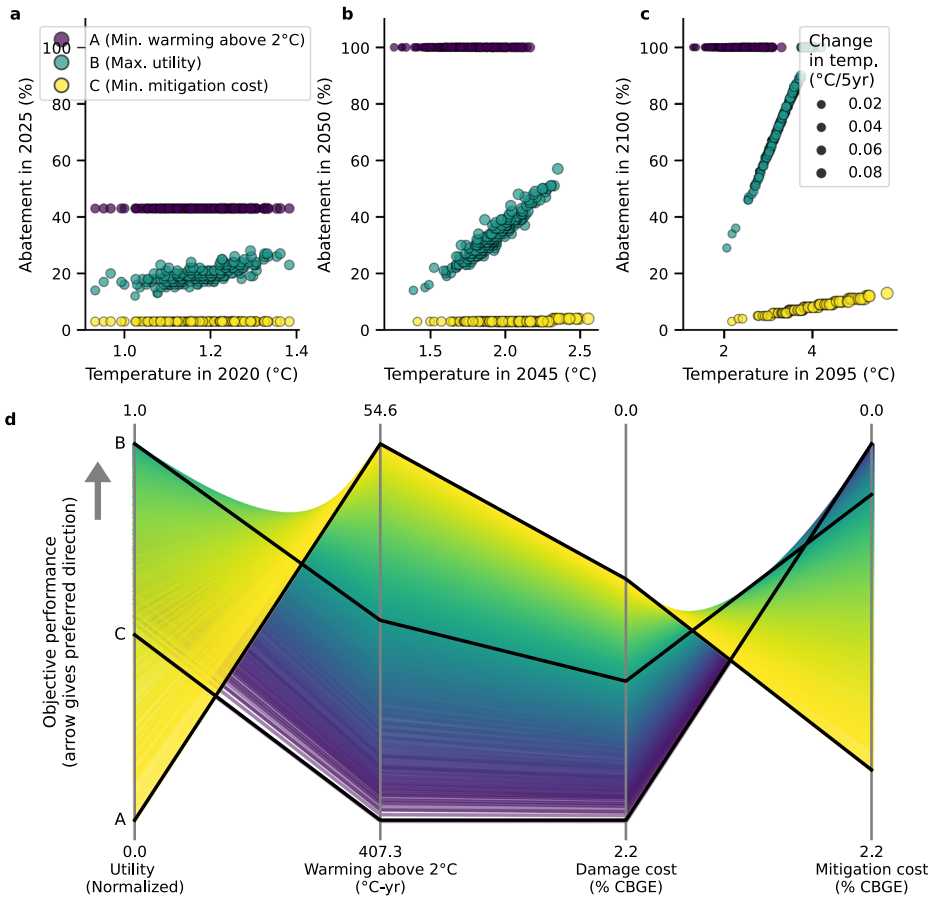


Fig. 1 Pareto-optimal adaptive strategies and resulting trade-offs between multiple mitigation objectives. **a–c** Abatement of CO₂ baseline emissions in one time step as a function of temperature level and rate of change. Each colored cloud of points is a different strategy. Strategy A (C) implies maximal (minimal) abatement irrespective of temperature detected. Compromise solutions like strategy B exhibit an intermediate behavior, increasing action with evidence of warming. **d** Parallel axes plot of Pareto-optimal solutions, sorted and colored by decreasing two-degree years. Strategies are horizontal lines crossing each vertical axis in correspondence to that axis dimension of performance. Strategy A (C) has minimal (maximal) warming and damages, and maximal (minimal) mitigation costs. Strategy B lies in between these extremes

pace of decarbonization. Aggressive early abatement efforts also drastically reduce climate change damages. Utility is lowest and mitigation costs are highest for this strategy: mitigation costs peak by 2050 and then decrease as technical progress drives down the costs of carbon-free energy (Figure S3). Even at maximum abatement, within inertia constraints, the system would experience an expected residual warming of 55 two-degree years with median temperatures peaking at 2.4 °C. This is far below the roughly 400 two-degree years that would be otherwise expected with business-as-usual emissions. Lower temperature outcomes require earlier availability of negative emissions, here assumed to be viable only after 2150 and limited to 20% of baseline emissions (see Supplementary Information for a discussion and sensitivity analysis on this assumption).

A solution that minimizes mitigation costs leaves emissions unabated (solution C in Fig. 1). This strategy represents a no-policy counterfactual scenario, commonly included for IAM scenario assessments (Riahi et al. 2017; Krey 2014) and often used in the expression of many Nationally Determined Contributions within the Paris Agreement (Vaidyula and Hood 2018). Such a baseline still accounts for an exogenous reduction in emission intensity of the economy, but has no explicit abatement. Our baseline emissions are within the uncertainty range assessed by the Shared Socio-Economic Pathways (Riahi et al. 2017) (Figure S13). Median temperatures are highest with this strategy, rising above 6 °C, causing the largest economic damages in the considered cases (Fig. 1d). Utility is partially penalized, mostly for future generations who will bear the consequences of this climate inaction. Abatement rises above zero by the end of this century, as fossil fuels become more scarce, damages increase, and discounted mitigation costs become negligible (Figure S3).

The full range of Pareto-optimal solutions reveals strong trade-offs between the considered objectives (Fig. 1d). Reducing climate damages increases mitigation costs and tightly follows a reduction in two-degree years. Utility decreases either when abatement costs or climate damages are high. Achieving maximum utility involves moderate mitigation, with an expected overshoot in temperature of around 200 two-degree years and a median temperature rise peaking at 3.8 °C (Figure S3).

The choice of preferred objective strongly influences the policy rule (Fig. 1a–c). Minimizing two-degree years requires the highest abatement, with no influence of observed temperatures (Fig. 1a–c). Abatement requirements are less stringent when maximizing utility: optimal behavior veers towards full abatement or inaction depending on revealed temperatures.

Adaptive strategies can provide considerable economic and environmental benefits (Fig. 2). We compare Pareto-optimal adaptive strategies with non-adaptive strategies (optimizing CO₂ reductions inter-temporally without feedbacks from the observed temperatures within each of the sampled SOWs). We obtain the inter-temporal strategies with the same evolutionary multi-objective optimization setup, where decisions correspond to the levels of abatement at each time step. Such non-adaptive strategies represent the framework most commonly used to analyze mitigation strategies (Clarke et al. 2014). We find that responding to observed warming improves the average performance compared to non-adaptive solutions. The greatest gains are attained relative to the trade-off between mitigation cost and two-degree years, representative of the tension between tolerable warming and abatement costs (Fig. 2a). For strategies at the extremes of the cost spectrum, action is largely independent of observations. In other words, full abatement and complete inaction are strategies that are virtually indistinguishable from their non-adaptive counterparts. Advantages of adaptive strategies increase when moving away from these extremes.

Adaptive strategies exploit temperature information feedbacks to reduce abatement costs and warming. This is possible because abatement increases and decreases in concert with the resulting temperature changes driven by increasing and decreasing climate sensitivities (Fig. 2b). For SOWs with low climate sensitivities, the benefits provided by the adaptive strategy are mostly economic. For cases with high climate sensitivities, the benefits due to the adaptive strategy are mostly in terms of avoided warming.

We illustrate how one adaptive solution (X in Fig. 2) improves on the performance of two non-adaptive solutions that share either the same two-degree years (Y in Fig. 2) or the same mitigation costs (Z in Fig. 2). For the same warming, the adaptive solution results in lower mitigation efforts (Fig. 2a and b). Exploiting the temperature information feedback, the peak in mitigation costs decreases by 0.1% of GDP loss with a probability of 50%, and by 0.7% of GDP loss with a probability of 25% (Figure S5b). For the same mitigation cost, the adaptive

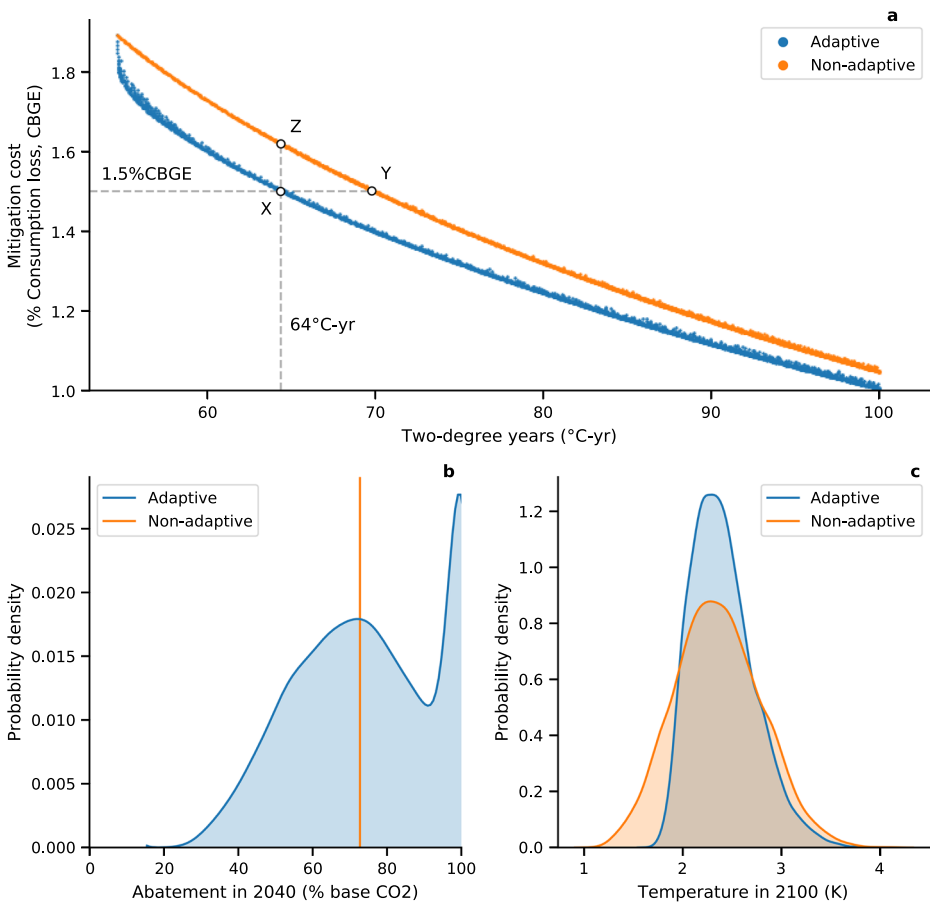


Fig. 2 Trade-off between mitigation cost and warming under adaptive and non-adaptive strategies. **a** Projection of the Pareto set of optimal solutions along the dimensions of two-degree years and mitigation cost outcomes, with some illustrative strategies highlighted. Strategy X is an adaptive solution compromising between economic performance and warming. Strategy Y is a non-adaptive solution with the same mitigation costs of X, and strategy Z is a non-adaptive solution with the same exceedance in two-degree years of X. **b** Distribution of abatement levels in 2040 for the two strategies with the same expected warming outcome. **c** Distribution of temperature increase in 2100 for the two strategies with the same expected mitigation costs

solution reduces the risk of high-temperature outcomes (Fig. 2a and c). Adaptively allocating the mitigation budget of X reduces the probability of exceeding 3 $^{\circ}\text{C}$ (4 $^{\circ}\text{C}$) by 25% (50%) (Figure S5c). These figures are illustrative and depend on the assumed distribution, as is often the case in probabilistic modelling. We guard against potential misrepresentations of reality by also considering alternative climate ensembles (Figure S12).

Considering deep uncertainty drastically increases the advantages of adaptive strategies. Two key factors affecting climate damages are subject to particularly divergent perspectives and assessments. One is climate sensitivity, along with the correlated quantities of ocean diffusivity and aerosol scaling (Knutti et al. 2017). The other is the shape and magnitude of the function relating economic losses to a given rise in temperatures (Weitzman 2012; Hwang et al. 2013; Goes et al. 2011).

We explore two additional ensembles of SOWs, representing a more sensitive and a less sensitive climate (Fig. 3b). These are consistent with historical data, and result from different choices of priors in the climate calibration. We also consider two additional shapes for the relationship between temperature increase and economic losses (Fig. 3a). These two representations of the damage relationship explore cases when warming is more damaging compared to the standard assumption in DICE (Nordhaus and Sztorc 2013). Expanding the scope of SOWs that capture the deeply uncertain nature of climate sensitivity and damages clarifies how the information feedback aids the ability of mitigation strategies to confront cases beyond those used in the initial optimization. To this end, we quantify the changes in the trade-off between mitigation costs and the 95th percentile of the damage costs, measured in Balanced-Growth Equivalent terms.

Re-evaluating strategies under these deeply uncertain damage and climate responses expands the range of economic losses. Damages now span two orders of magnitude of equivalent loss in today's consumption (Fig. 3c), ranging from 0.6% for high mitigation, low climate sensitivity and low damage response, to 40% for low mitigation, high climate sensitivity and high damage response. The trade-off curves between mitigation costs and worst-case damages exhibit similar features to those observed for the standard assumptions (corresponding to a relatively low damage function and a medium climate sensitivity).

Once committed to an expected mitigation cost budget, replacing a non-adaptive strategy with an adaptive one can yield large performance gains. For higher climate sensitivities and damages, the adaptive strategies can avoid losses as high as 10% of today's consumption compared to the non-adaptive counterparts (Fig. 3d). Information feedback is particularly beneficial when the mitigation cost budget is limited, as the strategy detects worst-case warming scenarios and allocates the scarce resources more efficiently. For lower climate sensitivities and damages, benefits are mostly in avoided mitigation efforts (green markers shifting to the left in Fig. 3c). Exposed to milder temperatures, adaptive strategies result in fewer expenditures than originally planned, with minor repercussions on extreme damage costs. For higher mitigation cost budgets, appreciable savings emerge both in damage and mitigation costs (oblique segments in Fig. 3c). Within each climate ensemble, adaptive strategies require fewer resources to respond with maximum abatement to high-impact scenarios. Overall, adaptive strategies ensure lower economic costs in a larger number of SOWs (Figure S11), while cutting the upper tail of future temperature increase (Figure S10).

How much adaptive strategies reduce extreme damage costs depends on both dimensions of deep uncertainty considered. A low climate sensitivity yields minimal increases in temperature, as well as damages for all considered damage functions. Hence, the benefits from adaptive abatement are relatively low. For higher temperature projections, the shape of the damage function matters considerably. The benefits of adaptive strategies are especially clear when damages quickly reach sizeable magnitudes as temperature rises, and when mitigation resources are scarce.

4 Conclusions

A predetermined mitigation strategy maximizing utility can offer a false sense of confidence in committing the world to a certain level of climate ambition. Here we showed how utility maximization, while compromising abatement and damage costs, and even if adapting to observed circumstances, skews preferences towards lower abatement costs and hide important trade-offs (Fig. 1). When giving more weight to a 2 °C target, adaptivity can reduce mitigation expenditures by 0.1–0.7% of GDP without worsening expected

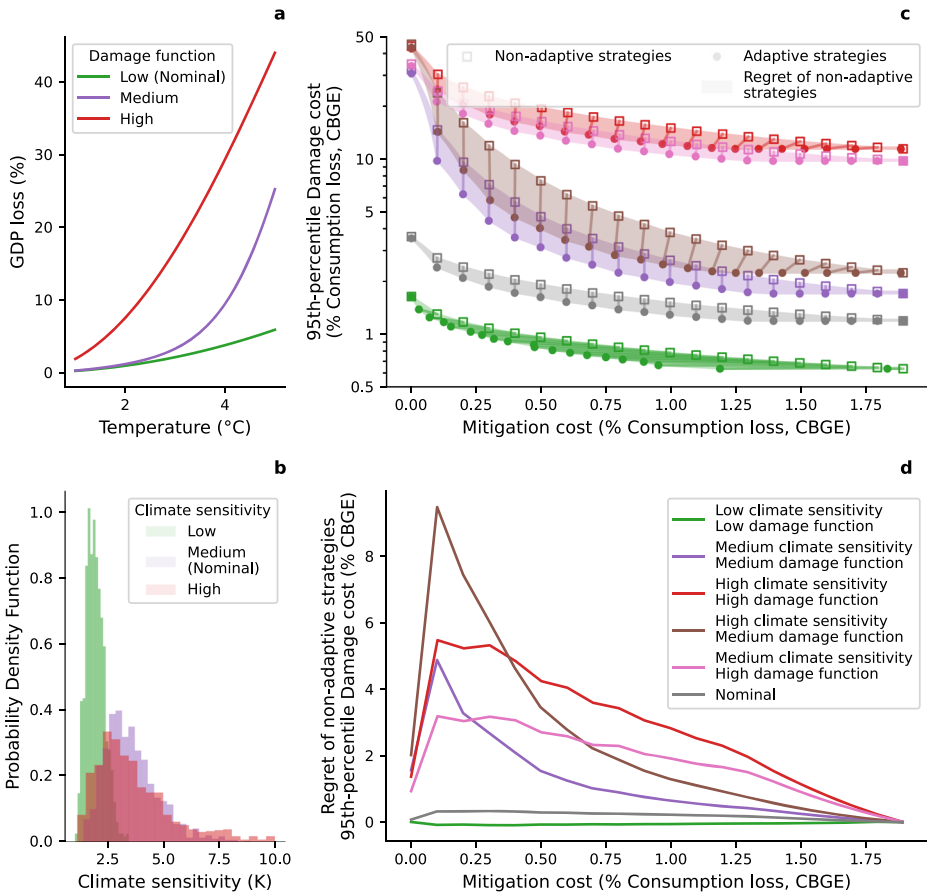


Fig. 3 Performance of adaptive vs non-adaptive strategies re-evaluated under alternative damage functions and climate responses. **a** Global GDP loss as a function of atmospheric temperature increase above pre-industrial levels for a low, medium and high impact case. **b** Posterior distributions of climate sensitivity for a low, medium and high sensitivity case. **c** 95th-percentile damage cost and mitigation costs of 20 non-adaptive (empty squares) and adaptive (filled circles) Pareto-optimal solutions. The 40 solutions are re-evaluated under alternative climate response and damage function assumptions, denoted by a different color. Shaded areas quantify the regret of not adapting. Segments connect solutions having the same mitigation costs under nominal conditions. The regret, in terms of reduced 95th percentile of damage costs, is shown explicitly in panel **d**. For visual clarity, only a subset of damage and climate combinations is shown

temperature overshoot, and halve the probability of exceeding a high temperature like 4 °C under high sensitivities for the same expected cost (Fig. 2). Beyond ensuring environmental and economic efficiency, adaptive strategies can avoid losses as high as 10% of today’s consumption compared to the non-adaptive counterparts when confronted with deeply uncertain climate responses (Fig. 3).

Adaptive strategies can reduce high regrets by timely adjusting action to extreme circumstances. They can also help to reduce policy deadlocks by transitioning the discussion to evidence-based indices of warming. While the global level of abatement is arguably hard to enforce in a top-down fashion and will likely remain an emergent property of individual state efforts, our adaptive planning framework can still be implemented by signalling

the need for a higher or lower mitigation response intensity at the global level. The recommended adaptive planning response could be negotiated through recurring international meetings (like the Conference of Parties) following the most recent scientific information. The corresponding signal needed for societal change could then be communicated through authoritative assessments (like the IPCC Assessment Reports), multi-party agreements (like the Paris Agreement), large-scale carbon pricing systems (like the EU Emissions Trading System), and shared financial instruments (like the Green Climate Fund). The benefit of our approach is that it illustrates the multi-faceted consequences of the available mitigation choices. This facilitates an iterative and collaborative process of convergence towards a shared level of climate action among stakeholders (Kwakkel and Haasnoot 2019).

The proposed approach lends itself to multiple extensions. For example, we adopted a simple and transparent model for conceptual clarity. Future research could use more complex IAMs to provide more detailed insights. Also, the framework is scalable to region-specific trajectories and objectives, to the extent it remains computationally feasible. Further analysis could include the exploration of other objectives (like health and energy security implications), and types of robust strategies (no-regret, risk-averse). Further research could look into the benefits of re-optimizing decisions at each time step based on an updated estimate of unknown parameters (model-predictive control), as well as incorporating additional estimators and signals to compensate for potential sources of stochasticity like natural temperature variability, measurement errors and shocks, which are not modeled here.

In our analysis, we explore the implication of global abatement aspiration levels, while leaving aside what deployment of policies is needed to achieve this aspiration level. Coupling this framework with more detailed models, and with the literature on climate policy instruments, could provide opportunities to better operationalize such aspirations. We also assume perfect observability of abatement levels and temperature changes, which might not be the case in reality, and further efforts should be allocated to ensure better monitoring for this approach to work. Another important channel for informing feedback policies could be the estimated entity of damage functions (Anthoff and Tol 2021): as more or less severe damage functions are ruled out or become less likely, damage-minimizing policies could follow as well.

The multi-objective framework adopted here opens new opportunities for including performance metrics otherwise hard to account for in traditional single-objective optimization modelling. Intergenerational and intragenerational justice, quantifiable with Gini or other inequality indices, could be directly included in the set of objectives, once the modelling is extended to a regional setup. This would help climate policy makers to immediately appreciate and compare equity implications of alternative strategies, alongside with the more traditionally discussed environmental and economic aspects.

Our approach provides a first set of benchmark strategies that policy-makers can already use to inform the design and implementation of climate mitigation solutions. These, at present, revolve mostly around prescribed CO₂ reduction schedules that maximize discounted expected utility, which are silent on the trade-offs between important objectives and the effects of deep and dynamic uncertainties. We demonstrate how moving to an adaptive strategy can drastically improve the ability to avoid high-regret futures as we navigate key abatement trade-offs in the face of these deep uncertainties.

Acknowledgements We thank I. Schaperdoth, R. Lempert, C. Helgeson, V. Srikrishnan, B. Lee, S. Wishbone, J. Oyler, C. Spence, K. Ruckert, T. Wong, M. Tavoni and the reviewers for their valuable inputs.

Author contribution G.M. and K.K. originally designed the study. All authors contributed to its final design. G.M. implemented the analysis. G.M. and K.K. wrote the first draft of the paper. All authors contributed to the final writing of the paper.

Funding Open access funding provided by Politecnico di Milano within the CRUI-CARE Agreement. This work was partially supported by the National Science Foundation through the Network for Sustainable Climate Risk Management (SCRiM) under NSF cooperative agreement GEO-1240507, the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013), the ERC grant agreement n° 336155 (COBHAM), the H2020 European Commission Project NAVIGATE under Grant Agreement n° 821124, and the Penn State Center for Climate Risk Management. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding entities.

Availability of data and material The dataset collected for the analysis is available upon reasonable request to the authors.

Code availability The code for the analysis and for generating all the figures is available online at <https://github.com/jackjackk/paper-dice-dps>.

Declarations

Conflict of interest The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.






References

- Anthoff D, Tol RSJ (2009) The impact of climate change on the balanced growth equivalent: An application of FUND. *Environ Resour Econ* 43(3):351–367. <https://doi.org/10.1007/s10640-009-9269-5>
- Anthoff D, Tol RSJ (2021) Testing the dismal theorem. <https://papers.ssrn.com/abstract=3807002>
- Bartholomew E, Kwakkel JH (2020) On considering robustness in the search phase of robust decision making: A comparison of many-objective robust decision making, multi-scenario many-objective robust decision making, and many objective robust optimization. *Environ Model Softw* 127:104699. <https://doi.org/10.1016/j.envsoft.2020.104699> <https://www.sciencedirect.com/science/article/pii/S136481521931223X>
- Burke M, Hsiang SM, Miguel E (2015) Global non-linear effect of temperature on economic production. *Nature* 527(7577):235–239. <https://doi.org/10.1038/nature15725>
- Cai Y, Lenton TM, Lontzek TS (2016) Risk of multiple interacting tipping points should encourage rapid CO₂ emission reduction. *Nat Climate Chang* 6(5):520–525. <https://doi.org/10.1038/nclimate2964>
- Cai Y, Lontzek T (2018) The social cost of carbon with economic and climate risks. <https://doi.org/10.1086/701890>
- Chylek P, Lohmann U (2008) Aerosol radiative forcing and climate sensitivity deduced from the Last Glacial Maximum to Holocene transition. *Geophys Res Lett* 35(4). <https://doi.org/10.1029/2007GL032759>
- Clarke L, Jiang K, Akimoto K, Babiker M, Blanford G, Fisher-Vanden K, Hourcade JC, Krey V, Kriegler E, Löschel A (2014) Assessing transformation pathways. In: Edenhofer O, Pichs-Madruga R, Sokona Y, Minx JC, Farahani E, Kadner S, Seyboth K, Adler A, Baum I, Brunner S, Eickemeier P, Kriemann B, Savolainen J, Schlömer S, von Stechow C, Zwickel T (eds) *Climate change 2014: Mitigation of climate change* vol Ch. 6. Cambridge University Press

- Fukuyama F (2007) Blindside: How to anticipate forcing events and wild cards in global politics. <http://www.jstor.org/stable/10.7864/j.ctt6wpff7>
- Fuss S, Canadell JG, Peters GP, Tavoni M, Andrew RM, Ciais P, Jackson RB, Jones CD, Kraxner F, Nakićenović N (2014) Betting on negative emissions. *Nat Clim Chang* 4(10):850
- Garner G, Reed P, Keller K (2016) Climate risk management requires explicit representation of societal trade-offs. *Clim Chang* 134(4):713–723. <https://doi.org/10.1007/s10584-016-1607-3>
- Giuliani M, Castelletti A, Pianosi F, Mason E, Reed PM (2015) Curses, tradeoffs, and scalable management: Advancing evolutionary multiobjective direct policy search to improve water reservoir operations. *J Water Resour Plan Manag* 142(2):04015050
- Goes M, Tuana N, Keller K (2011) The economics (or lack thereof) of aerosol geoengineering. *Clim Chang* 109(3-4):719–744. <https://doi.org/10.1007/s10584-010-9961-z>
- Grübler A, Nakićenović N, Victor DG (1999) Dynamics of energy technologies and global change. *Energy Policy* 27(5):247–280
- Hadka D, Reed P (2012) Borg: an auto-adaptive many-objective evolutionary computing framework. *Evol Comput* 21(2):231–259. https://doi.org/10.1162/EVCO_a_00075
- Haustein K, Allen MR, Forster PM, Otto FEL, Mitchell DM, Matthews HD, Frame DJ (2017) A real-time global warming index. *Sci Rep* 7(1):15417. <https://doi.org/10.1038/s41598-017-14828-5>
- Heal G (2017) The economics of the climate. *J Econ Lit* 55(3):1046–1063. <https://doi.org/10.1257/jel.20151335>
- Herman JD, Reed PM, Zeff HB, Characklis GW (2015) How should robustness be defined for water systems planning under change? *J Water Resour Plan Manag* 141(10):04015012. [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0000509](https://doi.org/10.1061/(ASCE)WR.1943-5452.0000509)
- Hobbs BF, Chankong V, Hamadeh W, Stakhiv EZ (1992) Does choice of multicriteria method matter? an experiment in water resources planning. *Water Resour Res* 28(7):1767–1779. <https://doi.org/10.1029/92WR00712>
- Hwang IC, Reynès F, Tol RSJ (2013) Climate policy under fat-tailed risk: an application of DICE. *Environ Resour Econ* 56(3):415–436. <https://doi.org/10.1007/s10640-013-9654-y>
- Jarvis AJ, Leedal DT, Hewitt CN (2012) Climate-society feedbacks and the avoidance of dangerous climate change. *Nat Clim Chang* 2(9):668–671. <https://doi.org/10.1038/nclimate1586>
- Jarvis A, Leedal D, Taylor CJ, Young P (2009) Stabilizing global mean surface temperature: A feedback control perspective. *Environ Model Softw* 24(5):665–674. <https://doi.org/10.1016/j.envsoft.2008.10.016> <http://www.sciencedirect.com/science/article/pii/S136481520800193X>
- Kasprzyk JR, Nataraj S, Reed PM, Lempert RJ (2013) Many objective robust decision making for complex environmental systems undergoing change. *Environ Model Softw* 42:55–71. <https://doi.org/10.1016/j.envsoft.2012.12.007> <https://www.sciencedirect.com/science/article/pii/S1364815212003131>
- Keller K, McInerney D, Bradford DF (2008) Carbon dioxide sequestration: how much and when? *Clim Chang* 88(3):267. <https://doi.org/10.1007/s10584-008-9417-x>
- Knutti R, Rugenstein MAA, Hegerl GC (2017) Beyond equilibrium climate sensitivity. *Nat Geosci* 10(10):727–736. <https://doi.org/10.1038/ngeo3017>
- Krey V (2014) Global energy-climate scenarios and models: a review. *WIREs Energy Environ* 3(4):363–383. <https://doi.org/10.1002/wene.98>
- Kriegler E (2005) Imprecise probability analysis for integrated assessment of climate change. PhD Thesis, Universität Potsdam
- Kunreuther H, Heal G, Allen M, Edenhofer O, Field CB, Yohe G (2013) Risk management and climate change. *Nat Clim Chang* 3(5):447–450. <https://doi.org/10.1038/nclimate1740>
- Kwakkel JH, Haasnoot M (2019) Supporting DMDU: A taxonomy of approaches and tools. In: Marchau VAWJ, Walker WE, Bloemen PJTM, Popper SW (eds) *Decision making under deep uncertainty: from theory to practice*. Springer International Publishing, pp 355–374. https://doi.org/10.1007/978-3-030-05252-2_15
- Kwakkel JH, Walker WE, Haasnoot M (2016) Coping with the wickedness of public policy problems: Approaches for decision making under deep uncertainty 142:(3):01816001. [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0000626](https://doi.org/10.1061/(ASCE)WR.1943-5452.0000626)
- Lamontagne JR, Reed PM, Marangoni G, Keller K, Garner GG (2019) Robust abatement pathways to tolerable climate futures require immediate global action. *Nat Clim Chang* 9(4):290. <https://doi.org/10.1038/s41558-019-0426-8>
- Lenton TM, Ciscar JC (2013) Integrating tipping points into climate impact assessments. *Clim Chang* 117(3):585–597. <https://doi.org/10.1007/s10584-012-0572-8>

- MacMartin DG, Kravitz B, Keith DW, Jarvis A (2014) Dynamics of the coupled human–climate system resulting from closed-loop control of solar geoengineering. *Clim Dynam* 43(1-2):243–258. <https://doi.org/10.1007/s00382-013-1822-9>
- McPhail C, Maier HR, Kwakkel JH, Giuliani M, Castelletti A, Westra S (2018) Robustness metrics: How are they calculated, when should they be used and why do they give different results? *Earth's future*, pp 169–191. [https://doi.org/10.1002/2017EF000649@10.1002/\(ISSN\)2328-4277.RESDEC1](https://doi.org/10.1002/2017EF000649@10.1002/(ISSN)2328-4277.RESDEC1) <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2017EF000649>
- Nordhaus WD (2017) Revisiting the social cost of carbon. *Proc Natl Acad Sci* 114(7):1518–1523. <https://doi.org/10.1073/pnas.1609244114> <http://www.pnas.org/content/114/7/1518>
- Nordhaus W, Sztorc P (2013) DICE 2013r: Introduction and user's manual. http://www.econ.yale.edu/nordhaus/homepage/homepage/documents/DICE_Manual_100413r1.pdf
- Pindyck RS (2013) Climate change policy: What do the models tell us? *J Econ Lit* 51(3):860–872. <https://doi.org/10.1257/jel.51.3.860>
- Riahi K, van Vuuren DP, Kriegler E, Edmonds J, O'Neill BC, Fujimori S, Bauer N, Calvin K, Dellink R, Fricko O, Lutz W, Popp A, Cuaresma JC, Kc S, Leimbach M, Jiang L, Kram T, Rao S, Emmerling J, Ebi K, Hasegawa T, Havlik P, Humpenöder F, Da Silva LA, Smith S, Stehfest E, Bosetti V, Eom J, Gernaat D, Masui T, Rogelj J, Strefler J, Drouet L, Krey V, Luderer G, Harmsen M, Takahashi K, Baumstark L, Doelman JC, Kainuma M, Klimont Z, Marangoni G, Lotze-Campen H, Obersteiner M, Tabeau A, Tavoni M (2017) The shared socioeconomic pathways and their energy, land use, and greenhouse gas emissions implications: An overview. *Glob Environ Chang* 42:153–168. <https://doi.org/10.1016/j.gloenvcha.2016.05.009> <http://www.sciencedirect.com/science/article/pii/S0959378016300681>
- Shayegh S, Thomas VM (2015) Adaptive stochastic integrated assessment modeling of optimal greenhouse gas emission reductions. *Clim Chang* 128(1-2):1–15. <https://doi.org/10.1007/s10584-014-1300-3>
- Stechow Cv, Minx JC, Riahi K, Jewell J, McCollum DL, Callaghan MW, Bertram C, Luderer G, Baiocchi G (2016) 2°C and SDGs: United they stand, divided they fall? *Environ Res Lett* 11(3):034022. <https://doi.org/10.1088/1748-9326/11/3/034022>
- UNFCCC (2015) Adoption of the Paris agreement. <http://unfccc.int/resource/docs/2015/cop21/eng/109r01.pdf>
- Urban NM, Holden PB, Edwards NR, Sriver RL, Keller K (2014) Historical and future learning about climate sensitivity. *Geophys Res Lett* 41(7):2543–2552
- Urban NM, Keller K (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(5):737–750. <http://onlinelibrary.wiley.com/doi/10.1111/j.1600-0870.2010.00471.x/full>
- Vaidyula M, Hood C (2018) Accounting for baseline targets in NDCs: Issues and options for guidance. <https://doi.org/10.1787/9ae65cc1-en>
- Vuuren DP, Hof AF, Sluisveld MA, Riahi K (2017) Open discussion of negative emissions is urgently needed. *Nat Energy* 2(12):902–904. <https://doi.org/10.1038/s41560-017-0055-2>
- Wagner G, Zeckhauser RJ (2018) Confronting deep and persistent climate uncertainty. <https://gwagner.com/wp-content/uploads/Wagner-Zeckhauser-181120-Deep-and-persistent-climate-uncertainty.pdf>
- Watson AA, Kasprzyk JR (2017) Incorporating deeply uncertain factors into the many objective search process. *Environ Model Softw* 89:159–171. <https://doi.org/10.1016/j.envsoft.2016.12.001>, <https://www.sciencedirect.com/science/article/pii/S1364815216310593>
- Weitzman ML (2011) Fat-tailed uncertainty in the economics of catastrophic climate change. *Rev Environ Econ Policy* 5(2):275–292. <https://doi.org/10.1093/reep/rer006> <https://academic.oup.com/reep/article/5/2/275/1565182>
- Weitzman ML (2012) GHG targets as insurance against catastrophic climate damages. *J Public Econ Theory* 14(2):221–244. <https://doi.org/10.1111/j.1467-9779.2011.01539.x>
- Wong TE, Bakker AMR, Ruckert K, Applegate P, Slangen ABA, Keller K (2017) BRICK v0.2, a simple, accessible, and transparent model framework for climate and regional sea-level projections. *Geosci Model Dev* 10(7):2741–2760. <https://www.geosci-model-dev.net/10/2741/2017/gmd-10-2741-2017.html>

Affiliations

Giacomo Marangoni^{1,2}  · **Jonathan R. Lamontagne**³  · **Julianne D. Quinn**⁴  · **Patrick M. Reed**⁵  · **Klaus Keller**⁶ 

¹ Department of Management, Economics and Industrial Engineering, Politecnico di Milano, Milan, Italy

² RFF-CMCC European Institute on Economics and the Environment, Centro Euro-Mediterraneo sui Cambiamenti Climatici, Milan, Italy

³ Department of Civil and Environmental Engineering, Tufts University, Medford, MA, USA

⁴ Department of Engineering Systems and Environment, University of Virginia, Charlottesville, VA, USA

⁵ School of Civil and Environmental Engineering, Cornell University, Ithaca, NY, USA

⁶ Department of Geosciences and Earth and Environmental Systems Institute, Penn State University, University Park, PA, USA