

Deep uncertainty quantification: with an application to integrated assessment models *

Preliminary: do not share without the consent of the authors

Aleksandra Friedl[†], Felix Kübler[‡], Simon Scheidegger[§], Takafumi Usui[¶]

First version: April 20, 2021

This version: November 30, 2023

Abstract

This paper presents a comprehensive method for efficiently solving stochastic Integrated Assessment Models (IAMs) and performing parametric uncertainty quantification. Our approach consists of two main components: a deep learning-based algorithm designed to globally solve IAMs as a function of endogenous and exogenous state variables as well as uncertain parameters within a single model evaluation. Additionally, we develop a Gaussian process-based surrogate model to facilitate the efficient analysis of key metrics, such as the social cost of carbon, with respect to uncertain model parameters. Our approach enables a rapid estimation of Sobol' indices, Shapley values, and univariate effects, which would otherwise be computationally very challenging. To demonstrate the effectiveness of our method, we posit a high-dimensional stochastic IAM that aligns with cutting-edge climate science. This model incorporates a social planner with recursive preferences, iterative belief updates of equilibrium climate sensitivity using Bayes' rule, and stochastic climate tipping. Our computations reveal that most of the variability in the social cost of carbon stems from the parametric uncertainty in the equilibrium climate sensitivity and in the damage function. We also show that the uncertainty about the equilibrium climate sensitivity resolves in about a decade, which in turn leads to higher optimal temperatures and a slightly decreased social cost of carbon compared to a modeling set-up without Bayesian learning.

Keywords: Climate policy, Epstein-Zin preferences, Climate tipping, Deep learning, Uncertainty quantification, Gaussian processes

JEL classification: C61, C63, D58, H23, Q54, Q58

*We thank Lars Hansen, Alena Miftakhova, Karl Schmedders, Christian Traeger, Frank Venmans, Pablo Winant, as well as seminar and conference participants at the University of Chicago, the University of Lausanne, Imperial College, ESSEC, CERGE-EI Prag, ARIC Hamburg, the University of Geneva, the University of Hamburg, the University of Colorado Boulder, the University of Pennsylvania, the Australian National University, SAET 2023, IE University, and the "online seminar series Applied Machine Learning, Economics, and Data Science (AMLEDS) for helpful comments and suggestions. This work was generously supported by the Swiss National Science Foundation (SNF), under project ID "Can Economic Policy Mitigate Climate-Change".

[†]Department of Economics, HEC Lausanne, University of Lausanne, Switzerland. Email: aleksandra.friedl@unil.ch

[‡]Department for Banking and Finance, University of Zürich, Switzerland, and SFI. Email: felix.kuebler@bf.uzh.ch

[§]Department of Economics, HEC Lausanne, University of Lausanne, Switzerland; Enterprise for Society (E4S). Email: simon.scheidegger@unil.ch

[¶]Department for Banking and Finance, University of Zürich, Switzerland. Email: takafumi.usui@bf.uzh.ch

1 Introduction

Integrated assessment models (IAMs) serve as a pivotal instrument for quantifying the prospective repercussions of anthropogenic climate change on mankind; they are a key factor in facilitating well-informed policy deliberations and in the prudent management of associated risks. Nevertheless, akin to the case with all economic models, the ramifications IAMs convey could occasionally exhibit considerable reliance on assumptions (which at times can be limiting) pertaining to the parameters and functional configurations embraced by the modeler. In light of this limitation, Pindyck (2013, p.860) issued a cautionary note regarding the broad application of IAMs, asserting that these models "*...have crucial flaws that make them close to useless as tools for policy analysis.*"

Rather than perceiving uncertainty as an insurmountable constraint on the applicability of this entire category of models, we embrace it by introducing a comprehensive framework for uncertainty quantification (UQ) within IAMs. Our approach consists of two parts: First, we present a stochastic IAM that entails Bayesian learning about the equilibrium climate sensitivity (ECS), aiming to capture probably the most prominent source of uncertainty inherent in IAMs.¹ Second, we devise a versatile and highly efficient computational technique named "Deep Uncertainty Quantification" (Deep UQ) to compute global solutions² and perform parametric UQ, which is particularly suitable for addressing the challenges posed by complex IAMs.

We posit a model that encompasses a comprehensive array of key components of IAMs. To the best of our knowledge, these components have not been collectively investigated in quantitative work due to the computational limitations of existing methods. Moreover, they have not been examined through the lens of UQ to study various quantities of interest (QoIs). These include the sensitivity of the social cost of carbon (SCC)³ concerning different types of uncertainty and the temporal evolution of its distribution. Our model is built on top of the influential DICE model (Nordhaus, 2017). However, we depart from it along several relevant dimensions. First, we align with recent advances in climate change economics by employing Epstein-Zin (-Weil) preferences (e.g., Epstein and Zin, 1989, Jensen and Traeger, 2014, Cai and Lontzek, 2019). Second, the climate physics module is designed in accordance with Folini et al. (2023) to ensure congruence with state-of-the-art climate science. One of the most iconic sources of uncertainty in climate science models is the ECS (Knutti et al., 2017), which is characterized by substantial uncertainty and fat-tailed probability distributions. Third, we consequently incorporate insights from the climate science literature (Roe and Baker, 2007) within our IAM framework and posit that the model's social planner can iteratively update her beliefs regarding ECS using Bayes' rule. This approach enables us to

¹ The ECS denotes the projected long-term temperature rise, that is, the equilibrium global mean near-surface air temperature, resulting from a doubling of atmospheric CO₂ concentration.

² We adopt the nomenclature from Brumm and Scheidegger (2017), referring to a "global solution" as a solution computed utilizing equilibrium conditions at numerous points within the state space of a dynamic model, as opposed to a "local solution" which relies on a local approximation around a model's steady state, as achieved through perturbation methods.

³ The SCC can be computed as the marginal cost of carbon emissions, which is the sum of all future damages resulting from an infinitesimal extra emission of CO₂ into the atmosphere, discounted at the market interest rate.

demonstrate the planner’s capacity to progressively diminish uncertainty surrounding the ECS over time. Fourth, to accurately represent the adverse economic consequences of anthropogenic climate change, we employ not only the damage function proposed by Nordhaus (2017) but also the model introduced by Weitzman (2012), which forecasts significantly more severe economic impacts as temperatures rise. Additionally, we enhance this model by integrating a stochastic tipping element, as suggested by Kotlikoff et al. (2021), to reflect the irreversible nature of climate change impacts, in line with the findings of Lenton et al. (2008).

To tackle the IAM mentioned above numerically, we introduce Deep UQ: a method that can efficiently solve nonlinear, non-stationary, high-dimensional dynamic stochastic climate economy models and perform UQ (Saltelli et al., 2007) at negligible computational costs, that is, within minutes to hours on a laptop rather than hundreds of thousands of hours on a supercomputer. An essential task in UQ, and the one we are concerned with in this paper, is to measure the relative importance of the various parameter inputs to QoIs, such as the SCC. Importance can be assessed via the effects of changing those inputs at random. This leads to a global sensitivity analysis (GSA; see, e.g., Saltelli et al. (2007)) in which statistical methods, based on an analysis of variance decomposition, measure variable importance. However, the measures traditionally used in the literature, that is, Sobol’ indices, univariate effects, and Shapley values (see, e.g., Owen (2014), Song et al. (2016), and references therein) typically require thousands, if not tens of thousands of model solutions to obtain convergent statistics (Harenberg et al., 2019). Thus, such an approach is a significant roadblock for performing UQ on stochastic IAMs, where a single model solution can take more than 100k CPU hours (Cai and Lontzek, 2019). Consequently, the previous literature on UQ in IAMs has typically been limited to local perturbations of solutions to stochastic models, or to GSA in non-stochastic settings (see, e.g., Anderson et al. (2014), Butler et al. (2014), Miftakhova (2021), and references therein).

Our proposed numerical solution framework for performing GSA in IAMs to alleviate the abovementioned challenges consists of two distinct parts. First, we enhance a generic, deep learning-based algorithm called “Deep Equilibrium Nets” (Azinovic et al., 2022) such that it can be used to solve stochastic IAMs globally as a function of the endogenous and exogenous state variables as well as their parameters at once in a single model solution. Such a “deep surrogate”—a high-precision approximation of an IAM based on deep neural networks⁴—greatly accelerates the model evaluations needed for UQ. This acceleration is achieved by reducing the computational load by orders of magnitude. Consequently, it enables a range of compute-intensive applications. Now, these tasks are simplified to interpolation tasks on the pre-computed optimal policies. These policies are functions of the state variables and parameters, rather than necessitating the repeated solution of an entire model from scratch for a fixed set of parameters.⁵ Next, we propose constructing a second, cheap-to-evaluate

⁴ For a thorough introduction to deep surrogates in economics and finance, see Chen et al. (2021), and references therein.

⁵ There are substantial computational bottlenecks in solving the type of IAMs described before as a function of endogenous and exogenous states as well as parameters in a single model evaluation globally because of i) the presence of random shocks, ii) a high-dimensional state space, iii) strong nonlinearities in the optimal policies (introduced, for instance, by the presence of highly non-linear damage

surrogate model for the QoIs, such as SCC, to perform GSA by using the former neural network surrogate of the IAM as an input. For computational tractability, we suggest using Gaussian processes (GPs; see, e.g., [Renner and Scheidegger \(2018\)](#), and references therein) as the surrogate modeling technique to approximate QoIs as a function of the model parameters.⁶

Applying our Deep UQ framework to a stochastic IAM with and without Bayesian learning about the ECS reveals several notable findings. By comparing the optimal policies from the IAMs with and without learning, we find that in settings without Bayesian learning, the agent opts for a SCC which is substantially higher than in an IAM with learning to ensure the abatement against possible high-temperature shocks, regardless of whether they occur or not. In contrast, when we consider an IAM with learning, it is optimal for the agent to allow for more emissions to learn about the ECS. This, in turn, leads to a realized SCC, which, in some cases, only reaches about half the level of the models with no learning. Next, we apply our UQ apparatus to analyze the impact of the uncertain parameters on the SCC through the lens of the Sobol’ indices, Shapley values, and univariate effects. Our numerical results show that the uncertainty stemming from the ECS and the one from the parameterization of the damage function are the two sources driving the variance of the SCC. In addition, our numerical experiments show that the uncertainty about the ECS resolves in about a decade, which in turn leads to higher optimal temperatures and a substantially increased SCC compared to the modeling set-up without Bayesian learning. Moreover, due to our Deep UQ methodology, all these results could be obtained by computationally cheap interpolations on a surrogate rather than repeatedly recalculating and simulating expensive models for a vast set of parameters, thereby making large-scale IAMs readily available to researchers merely having access to laptops rather than supercomputers.

The remainder of this article is organized as follows. In Section 2, we provide a brief review of the related literature. In Section 3, we posit a stochastic IAM with Bayesian learning about the ECS. Section 4 introduces our generic Deep UQ framework, and discusses its computational advantages relative to other existing methods. Section 5 presents our numerical experiments and results. Section 6 finally concludes.

functions (see, e.g., [Cai and Lontzek \(2019\)](#), and references therein), iv) irregular, that is, non-hypercubic geometries of the set of states visited along a simulation, and v) non-stationarity. In the presence of these five features, the curse of dimensionality ([Bellman, 1961](#)) imposes a considerable roadblock. While some methods can handle a subset of i) - v), most fail at matching all five requirements. Currently, and to the best of our knowledge, the only numerically tractable method that jointly addresses all five features relies on deep neural networks.

⁶ In this study, we deliberately construct the surrogate models using two different function approximators. As a rule of thumb, one should use a surrogate model based on deep neural networks when dealing with large datasets that require learning complex nonlinear relationships. For example, in the context of solving our stochastic, nonlinear IAMs, we generate billions of observations in the solution process to train our neural network. Hence, we apply neural networks. Conversely, surrogate models based on GPs are appropriate when dealing with small datasets or when interpretability is a priority over predictive accuracy. Consequently, since the SCC and other QoIs are based on computationally relatively expensive simulations based on the model solution in our concrete case, we apply GPs in the latter case.

2 Literature review

This paper is related to four strands of the literature: i) research focusing on stochastic IAMs; ii) studies investigating the uncertain ECS, such as through Bayesian learning; iii) global solution methods based on machine learning; and iv) parametric UQ in IAMs.⁷

First, our paper contributes to a fast-growing body of literature on endogenizing economic and climate uncertainty in IAMs to investigate, among other quantities, the SCC. Various studies (see, e.g., [Golosov et al. \(2014\)](#), [Traeger \(2021\)](#)) solve stochastic IAMs analytically and offer closed-form representations of the SCC. However, the elegance of these solutions comes at the expense of strong assumptions on the carbon cycle and other functional forms that enter the model. Tackling IAMs with uncertainty while trying to maintain numerical tractability by introducing various assumptions, [Traeger \(2014\)](#) implements a state-reduced and annually calibrated recursive dynamic programming version of the canonical DICE-2007 model ([Nordhaus, 2008](#)), to study uncertainty in the damage function ([Croston and Traeger, 2013, 2014](#)), to examine long-run growth risk ([Jensen and Traeger, 2014](#)), and to introduce climate tipping points ([Lemoine and Traeger, 2014, 2016](#)). In contrast, there is a relatively large body of literature that is based on deterministic finite-horizon models that investigate the implications of uncertainty by Monte-Carlo sampling in the parameter space (see, e.g., [Anthoff et al. \(2009\)](#), [Ackerman et al. \(2010\)](#), [Gillingham et al. \(2018\)](#), [Nordhaus \(2018\)](#)). One of the disadvantages of the latter approach is that “...this first-order approximation to stochastic analysis does not model a decision maker’s optimal response to uncertainty...”, as [Jensen and Traeger \(2014\)](#) point out. In a series of papers, [Cai and Lontzek \(2019\)](#), [Lontzek and Narita \(2011\)](#), [Lontzek et al. \(2015\)](#), [Cai et al. \(2015, 2016\)](#), [Lontzek et al. \(2016\)](#) examine the SCC under long-run growth risk and irreversible climate tipping, dropping many analytical simplifications and resorting to high-performance computing. [Barnett \(2023\)](#) applies an asset pricing approach to assess uncertainty related to climate change impact, and [van der Ploeg \(2021\)](#) takes uncertainty into account, deriving the risk-adjusted discount rate. The work presented here differs from previous studies in that we introduce a deep learning-based global solution method to solve stochastic IAMs of unprecedented complexity without needing to analytically simplify IAMs in a restrictive way or to tap into high-performance computing.

Second, our work is related to studies examining the interplay of uncertainty in the ECS ([Knutti et al., 2017](#)) and the economy. [Kelly and Kolstad \(1999\)](#), [Leach \(2007\)](#), [Jensen and Traeger \(2013\)](#), [Kelly and Tan \(2015\)](#), [Fitzpatrick and Kelly \(2017\)](#), [Hwang et al. \(2017\)](#), developed dynamic stochastic IAMs and considered Bayesian learning to update a posterior distribution of ECS, which is unknown to the social planner. The literature found that the social planner with Bayesian learning resolves fat-tailed uncertainty in ECS in the long run. We build on the literature of Bayesian learning over the ECS, and are, to the best of our knowledge, the first to combine it with Epstein-Zin preferences to capture the risk preferences and intertemporal consumption decisions.

Third, our work contributes to the emerging but fast-growing literature on machine learning-based solution methods in dynamic economic models (see, e.g., [Azinovic et al. \(2022\)](#), [Maliar et al. \(2021\)](#), [Ebrahimi Kahou et al. \(2021\)](#), [Fernández-Villaverde](#)

⁷ [Hassler et al. \(2016\)](#) provide an excellent general introduction and review on IAMs.

et al. (2023), Renner and Scheidegger (2018), Han et al. (2021), Barnett et al. (2023), Azinovic and Žemlička (2023)). The numerical solution to stochastic IAMs is often achieved by applying traditional, grid-based value function iteration (see, e.g., Cai and Lontzek (2019), and references therein). However, while highly successful in small to mid-scale models, such methods are strongly limited when the model complexity increases and, consequently, often need to resort to high-performance computers. In our work, we adopt the “Deep Equilibrium Nets” algorithm (Azinovic et al., 2022, Folini et al., 2023) to the context of stochastic non-stationary IAMs. To the best of our knowledge, we are the first to implement a deep learning-based solution method to solve a high-dimensional stochastic IAM with Bayesian learning about the equilibrium climate sensitivity parameter and to propagate parametric uncertainty using a Gaussian process-based surrogate model when estimating the social cost of carbon in a computationally efficient fashion; all our numerical results can be computed on a laptop. Consequently, we hope that this methodological contribution will enable the IAM community to tackle much richer models without needing to use hundreds of thousands of CPU hours of computing time on a supercomputer to solve a single model specification.

Fourth, our work contributes to studies that examine parametric uncertainty in IAMs.⁸ The IAM literature typically performs either local sensitivity analysis (see, e.g., Nordhaus (2008), Ackerman et al. (2010)) or GSA to study parametric uncertainty. Local sensitivity analysis is a one-at-a-time approach (Saltelli et al., 2007), where only one parameter changes values, keeping the other parameters fixed and analyzing the effect of the parameter on a quantity of interest such as the SCC. However, Saltelli and D’Hombres (2010), among others, contended that this approach tends to be unstructured and, more importantly, suffers from the fact that it is only local, that is, highly dependent on the chosen parameter values. Moreover, this approach cannot account for possible interactions between parameters and nonlinear relationships that are often encountered in economic models, thereby failing to offer reliable policy recommendations (Croston and Traeger, 2013). In contrast, GSA can alleviate these shortcomings (see, e.g., Smith (2014) for a review). In the GSA literature, the so-called Sobol’ indices, as well as univariate effects, are standard metrics to measure the importance of some input parameters to overall model outcome uncertainty (Harenberg et al., 2019). The Shapley value is an alternative, important metric in the context of UQ but which, to the best of our knowledge, has not yet been used in the context of IAMs. Owen (2014) adopted the idea of Shapley values from the cooperative game theory literature and demonstrated that the Shapley values intuitively attribute the overall variance of the model outcome to some input variables but not necessarily to be matched to the Sobol’ indices. The GSA approach is usually based on representing the uncertainty about each parameter by a (potentially bounded) probability distribution. This parameter uncertainty is propagated through the economic model by the repeated evaluation

⁸ Other sources of uncertainty, which we do not pursue in this paper but which have been studied in the literature, include model misspecification and ambiguity (see, e.g., Barnett et al. (2020), Zhao et al. (2023)). Other sources of uncertainty encompass climate tipping points (Lenton et al., 2008), the damage function, and the transition to green technology. Weitzman (2012), for instance, introduced a tipping point to the standard Nordhaus damage function (Nordhaus, 2017), which dramatically raises damages for temperature increases beyond a given level.

of randomly drawn parameter vectors. The required sampling from the parameter distributions is usually done by Monte Carlo simulation. However, due to the slow convergence properties of Monte Carlo simulations, a very large number of draws is usually required, which, prior to our work, was a prohibitive roadblock to performing GSA with globally solved stochastic IAMs. [Anderson et al. \(2014\)](#) and [Butler et al. \(2014\)](#) are early examples where GSA is applied to the canonical, non-stochastic DICE model based on a large number of Monte-Carlo samples. More recently, [Miftakhova \(2021\)](#) constructed surrogate models of the deterministic DICE model using the polynomial chaos expansion method and analyzed the Sobol’ indices and univariate effects of the SCC on various input parameters.⁹ We overcome these limitations by proposing to use neural network-based ([Chen et al., 2021](#)) as well as Gaussian process-based surrogate models ([Scheidegger and Bilionis, 2019](#)) for the IAM and the related QoIs such as the SCC. This measure significantly accelerates computations by orders of magnitude, transforming otherwise computationally expensive model solutions and simulations into straightforward interpolation tasks. This advancement enables, in this paper, to the best of our knowledge, the first application of GSA in stochastic IAMs with Bayesian learning about the ECS.

3 A stochastic IAM with Bayesian learning

In this section, we posit a discrete-time stochastic IAM with Bayesian learning about the uncertain ECS. Our model combines a stochastic representative agent framework with Epstein-Zin (-Weil) preferences ([Epstein and Zin, 1989](#), [Weil, 1989](#)), a climate emulator¹⁰ that follows the specification of [Folini et al. \(2023\)](#), and stochastic tipping points modeled in the spirit of [Kotlikoff et al. \(2021\)](#).

We proceed as follows: First, in Section 3.1, we provide a brief summary of the economic building blocks of the model. Second, in Section 3.2, we specify the climate externality used in this work. Third, Section 3.3 describes our approach to modeling the learning process concerning the ECS. Finally, in Section 3.4, we present the recursive formulation of our model. In addition, further details about the model, such as calibration details, are provided in Appendix A.

3.1 The economic model

The economic side of our IAM builds on [Nordhaus \(2017\)](#), that is, a growth model that assumes a single, infinitely-lived representative agent, there are greenhouse gas (GHG) emissions due to production activity, and productivity is impacted by the state of the climate via a damage function.

We adopt the Epstein-Zin (-Weil) preferences ([Epstein and Zin, 1989](#), [Weil, 1989](#)) to address the critical aspects of risk preferences and intertemporal consumption de-

⁹ [Harenberg et al. \(2019\)](#) is the first example in economics to present GSA based on a polynomial chaos expansion of a standard real-business-cycle model.

¹⁰ To analyze climate change mitigation strategies, economists rely on simplified climate models – so-called climate emulators – that provide a realistic quantitative link between CO₂ emissions and global warming at low computational costs.

cisions in long-run climate policy discussions (see, e.g., [Crost and Traeger \(2013\)](#), [Cai and Lontzek \(2019\)](#)). The social welfare function U_t is recursively defined as follows:

$$U_t = \left[(1 - \beta) \frac{(C_t/L_t)^{1-1/\psi}}{1 - 1/\psi} L_t + e^{-\rho} \mathbb{E}_t \left[U_{t+1}^{1-\gamma} \right]^{\frac{1-1/\psi}{1-\gamma}} \right]^{\frac{1}{1-1/\psi}}, \quad (1)$$

where time t is discrete and measured in years, C_t represents consumption, and L_t denotes an exogenous population path adopted from [Nordhaus \(2017\)](#). Furthermore, ρ denotes the pure rate of time preference, ψ measures the intertemporal elasticity of substitution, γ represents the Arrow-Pratt risk aversion parameter, and $\mathbb{E}_t[\cdot]$ denotes the expectation operator conditional on time t .

The gross output Y_t^{Gross} in the economy is given by a Cobb-Douglas production function with labor-augmenting technological progress, that is,

$$Y_t^{\text{Gross}} = K_t^\alpha (A_t L_t)^{1-\alpha}, \quad (2)$$

where α represents the elasticity of substitution of capital, and A_t is a deterministic trend. There is an externality to production in the form of industrial carbon emissions, $E_{\text{Ind},t}$, that reads as

$$E_{\text{Ind},t} = \sigma_t (1 - \mu_t) Y_t^{\text{Gross}}. \quad (3)$$

Through the emission flow, the economic side of our IAM affects the climate (cf. Section 3.2 below). We follow [Nordhaus \(2017\)](#) and assume that the emissions produced are proportional to the final output, and scaled with an exogenous emission intensity σ_t . In our model, the emissions can be abated, that is, reduced, at a rate $\mu_t \in [0, 1]$. A zero abatement rate corresponds to the business-as-usual (BAU) case, whereas a value of $\mu = 1$, abatement corresponds to full mitigation. As in [Nordhaus \(2017\)](#), we assume that abatement is costly for the social planner and depends on the abatement rate with the following functional form:

$$\Theta_t(\mu_t) = \theta_{1,t} \mu_t^{\theta_2}. \quad (4)$$

In Eq. (4), $\theta_{1,t}$ is an exogenous process for the evolution of the abatement cost, and θ_2 is a parameter of the abatement cost function.

The unabated emissions enter the climate system and ultimately lead to an increase in the atmospheric temperature $T_{\text{AT},t}$ (cf. Section 3.2). A rise in the atmospheric temperature triggers damages that, in turn, will affect the final output. In our modeling setup, we employ two types of damage function commonly used in the IAM literature. The first one, denoted by $\Omega_t^N(T_{\text{AT},t})$, is the quadratic damage function by [Nordhaus \(2017\)](#):

$$\Omega_t^N(T_{\text{AT},t}) = \frac{1}{1 + \pi_1 T_{\text{AT},t} + \pi_2 T_{\text{AT},t}^2}. \quad (5)$$

We also employ the [Weitzman \(2012\)](#) damage function, denoted by $\Omega_t(T_{\text{AT},t})$, but

enhance it with stochastic tipping:¹¹

$$\Omega_t(T_{AT,t}) = \frac{1}{1 + \left(\frac{1}{\psi_1} T_{AT,t}\right)^2 + \left(\frac{1}{2 \cdot TP_t} T_{AT,t}\right)^{6.754}}. \quad (6)$$

We follow [Kotlikoff et al. \(2021\)](#) and assume that climate tipping is denoted by the variable TP_t that follows a random walk¹² with the Gaussian innovations with zero mean and variance equal to S_{TP} :¹³

$$TP_{t+1} = TP_t + \epsilon_{TP,t+1}, \quad \epsilon_{TP,t} \sim \mathcal{N}(0, S_{TP}). \quad (7)$$

The abatement costs and the damages jointly reduce the gross output of the economy, leading to the following budget constraint:

$$(1 - \Theta_t(\mu_t)) \Omega_t(T_{AT,t}) Y_t^{\text{Gross}} - C_t - I_t = 0. \quad (8)$$

The capital stock evolves according to

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

where δ is the depreciation rate. More details on the exogenous variables and the calibration are provided in Appendix A.1.

3.2 The climate model

We employ a climate emulator following the [Folini et al. \(2023\)](#) specification. This climate module builds on [Nordhaus \(2017\)](#) and comprises two fundamental building blocks: i) three stacked carbon reservoirs that represent the global carbon cycle and ii) a two-layer energy balance model.

There are two sources of emissions that enter the simple carbon cycle: industrial emissions $E_{\text{Ind},t}$, as defined in Eq. (3), which are related to the economic activity of the representative agent, and exogenous emissions $E_{\text{Land},t}$. Their sum $E_t = E_{\text{Ind},t} + E_{\text{Land},t}$ constitutes the total emission flow entering the carbon cycle.

The climate emulator consists of three carbon reservoirs that are represented by a three-dimensional vector $\mathbb{M}_t = (M_{AT,t}, M_{\text{UO},t}, M_{\text{LO},t})$, where the entries represent the mass of carbon in the atmosphere (AT), upper ocean (UO) and lower ocean (LO) respectively. The concentration of carbon in these three reservoirs evolves according

¹¹ Climate tipping is the occurrence of climate events that lead to irreversible damages to the environment. For more details, see, e.g., [Lenton et al. \(2008\)](#), and references therein.

¹² The original work by [Weitzman \(2012\)](#) calibrates $2 \cdot TP_t = 6.081$ to be constant over time. The climate tipping event occurs at about 3 °C excess temperature in the atmosphere. While the temperature change is below 3 °C, the climate tipping damage is negligibly small in the damage function. However, when it exceeds around 3 °C, the climate tipping damage begins to dominate the damage function. The way we model the tipping in a stochastic way reflects the fact that very little is known about the exact level of temperatures at which tipping could occur (see [Cai and Lontzek \(2019\)](#), and references therein).

¹³ In our computations below, we will truncate the normal distribution given in Eq. (7) to ensure numerical tractability (cf. Appendix A.2).

to a diffusion process described by the following expression:

$$\mathbb{M}_{t+1} = B\mathbb{M}_{t+1} + E_t, \quad (10)$$

where B is given by

$$B = \begin{pmatrix} 1 - b_{12} & b_{21} & 0 \\ b_{12} & 1 - b_{21} - b_{23} & b_{32} \\ 0 & b_{23} & 1 - b_{32} \end{pmatrix}. \quad (11)$$

The diffusion coefficients $b_{12}, b_{21}, b_{23}, b_{32}$ that we use in our computations are taken from [Folini et al. \(2023\)](#).

The baseline energy balance system, which we will have to slightly modify in Section 3.3 due to Bayesian learning, comprises two layers: the temperature of the atmosphere, $T_{AT,t}$ (representing the atmosphere and upper ocean), and the temperature of the ocean, $T_{OC,t}$. These layers are described by the following equations:

$$T_{AT,t+1} = \left(1 - c_1 c_3 - c_1 \frac{F_{2xco2}}{\Delta T_{AT,\times 2}}\right) T_{AT,t} + c_1 c_3 T_{OC,t} + c_1 \left(F_{2xco2} \log_2 \left(\frac{M_{AT,t}}{M_{AT}^*}\right) + F_{EX,t}\right), \quad (12)$$

$$T_{OC,t+1} = c_4 T_{AT,t} + (1 - c_4) T_{OC,t}. \quad (13)$$

F_{2xco2} denotes the forcing of equilibrium CO₂ doubling, $\Delta T_{AT,\times 2}$ is the ECS, M_{AT}^* stands for the equilibrium mass of carbon in the atmosphere, $F_{EX,t}$ represents exogenous radiative forcing, and c_1, c_3, c_4 are temperature-related parameters. All values, except the ECS, are calibrated according to [Folini et al. \(2023\)](#). A detailed discussion about the ECS follows in the subsequent Section 3.3. More details on the calibration of our climate emulator are given in Appendix A.2.

3.3 Bayesian learning over the equilibrium climate sensitivity

The response of the climate to emissions in the form of an increase in atmospheric temperature is subject to uncertainty in many ways. However, the ECS is presumably the most controversial number that determines how severe climate change will be in the long run ([Knutti et al., 2017](#)), as arguments between [Roe and Baker \(2007\)](#) and [Zaliapin and Ghil \(2010\)](#) in the climate science community illustrate. Therefore, to take this debate seriously in our IAM, Eq. (12) has to be modified appropriately.

Empirical observations, as well as studies based on energy-balanced models and global climate models, suggest that the probability density function of the ECS is most likely to peak in the region [2.0°C, 4.5°C], but this density exhibits a highly skewed, fat-tailed distribution ([Roe and Baker, 2007](#)). These authors, among others, pointed out that the ECS relates to the following feedback process:

$$T_{AT,\times 2} = \frac{T_{AT,\times 2}^0}{1 - f}, \quad (14)$$

where f represents the so-called climate feedback parameter, and $T_{AT,\times 2}^0$ is a reference climate sensitivity parameter, and which is assumed to be known to the social planner. The uncertainty in the ECS can be represented by the uncertainty in the climate feedback parameter f . Following [Roe and Baker \(2007\)](#), we assume that f follows a Gaussian distribution. By substituting Eq. (14) into Eq. (12), we obtain

$$T_{AT,t+1} = \left(c_1 \frac{F_{2\times\text{co2}}}{\Delta T_{AT,\times 2}^0} f + 1 - c_1 c_3 - c_1 \frac{F_{2\times\text{co2}}}{\Delta T_{AT,\times 2}^0} \right) T_{AT,t} + c_1 c_3 T_{OC,t} + c_1 \left(F_{2\times\text{co2}} \log_2 \left(\frac{M_{AT,t}}{M_{AT}^*} \right) + F_{EX,t} \right) + \epsilon_{T,t+1}, \quad (15)$$

where the uncertainty about the ECS is now modeled indirectly via f , and $\epsilon_{T,t+1}$ is a stochastic weather shock (cf. [Kelly and Tan \(2015\)](#)) which the social planner cannot observe.

In line with previous studies on Bayesian learning in IAMs (see, e.g., [Kelly and Kolstad \(1999\)](#), [Leach \(2007\)](#), [Webster et al. \(2008\)](#), [Kelly and Tan \(2015\)](#), [Hwang et al. \(2017\)](#), [Fitzpatrick and Kelly \(2017\)](#)), we update the social planner's belief on the climate feedback parameter f , and not the ECS, which is uncertain to the social planner. The social planner subjectively assumes her belief about the climate feedback parameter in period t , that is, \tilde{f}_{t+1} , when solving the dynamic programming problem. As [Roe and Baker \(2007\)](#) suggest,¹⁴ we assume that her prior belief on \tilde{f}_{t+1} follows a Gaussian distribution with mean $\mu_{f,t}$ and variance $S_{f,t}$. We also assume that the distribution is truncated from below at $\underline{f} = 0.4$ and from above at $\bar{f} = 0.9$. This implies that $\tilde{f}_{t+1} \sim \mathcal{N}(\mu_{f,t}, S_{f,t}, \underline{f}, \bar{f})$, and leads to an enhanced expression for the evolution of the atmospheric temperature, that is,

$$T_{AT,t+1} = \left(c_1 \frac{F_{2\times\text{co2}}}{\Delta T_{AT,\times 2}^0} \tilde{f}_{t+1} + 1 - c_1 c_3 - c_1 \frac{F_{2\times\text{co2}}}{\Delta T_{AT,\times 2}^0} \right) T_{AT,t} + c_1 c_3 T_{OC,t} + c_1 \left(F_{2\times\text{co2}} \log_2 \left(\frac{M_{AT,t}}{M_{AT}^*} \right) + F_{EX,t} \right) + \tilde{\epsilon}_{T,t+1}. \quad (16)$$

In period t , we denote the prior probability function of the social planner's belief of the climate feedback parameter, which follows the Gaussian distribution with mean μ_t^f and variance $S_{f,t}$, as $p(f)$. Given the prior $p(f)$, the planner observes the atmospheric temperature change and the likelihood function, which is also Gaussian, that is, $p(T_{AT} | f)$. By applying Bayes' rule, the posterior probability function of the climate

¹⁴ We choose the initial values for the prior mean and prior variance according to [Roe and Baker \(2007\)](#). We set the reference climate sensitivity as $\Delta T_{AT,\times 2}^0 = 1.2^\circ\text{C}$. Given the reference climate sensitivity, the initial climate sensitivity in the problem is given by $\frac{\Delta T_{AT,\times 2}^0}{1-f} = \frac{1.2}{1-0.65} = 3.42$. The latter value is not far from that of [Folini et al. \(2023\)](#) who determine the ECS as $\Delta T_{AT,\times 2} = 3.25^\circ\text{C}$. Thus, we modify no other parameters of the energy balance equations in order to be in line with the ECS value, but keep the remaining parameters at the values proposed by [Folini et al. \(2023\)](#).

feedback parameter given the observation T_{AT} is also Gaussian distributed and reads as:

$$p(f | T_{AT}) \propto p(T_{AT} | f) \times p(f). \quad (17)$$

As [Kelly and Tan \(2015\)](#), we rewrite the two stochastic components in Eq. (16) as

$$\begin{aligned} \left(1 - c_1 c_3 - c_1 \frac{F_{2\text{xco}_2}}{\Delta T_{AT,\text{x}2}^0}\right) T_{AT,t} + c_1 c_3 T_{OC,t} + c_1 \left(F_{2\text{xco}_2} \log_2 \left(\frac{M_{AT,t}}{M_{AT}^*}\right) + F_{EX,t}\right) \\ + \varphi_{1C} \tilde{f}_{t+1} T_{AT,t} + \tilde{\epsilon}_{T,t+1} - T_{AT,t+1} = 0, \end{aligned} \quad (18)$$

where $\varphi_{1C} = c_1 \frac{F_{2\text{xco}_2}}{\Delta T_{AT,\text{x}2}^0}$, and where $\tilde{f}_{t+1} \sim \mathcal{N}(\mu_{f,t}, S_{f,t}, \underline{f}, \bar{f})$, $\tilde{\epsilon}_{T,t+1} \sim \mathcal{N}(0, S_{\epsilon_T})$. Finally, by applying [DeGroot \(1970, p. 167, Theorem 1\)](#), we can analytically compute the posterior mean and the variance, that is,

$$\mu_{f,t+1} = \frac{S_{\epsilon_T} \mu_{f,t} + \varphi_{1C} T_{AT,t} \left(\varphi_{1C} T_{AT,t} \tilde{f}_{t+1} + \tilde{\epsilon}_{T,t+1}\right) S_{f,t}}{S_{\epsilon_T} + (\varphi_{1C} T_{AT,t})^2 S_{f,t}}, \quad (19)$$

$$S_{f,t+1} = \frac{S_{\epsilon_T} S_{f,t}}{S_{\epsilon_T} + (\varphi_{1C} T_{AT,t})^2 S_{f,t}}. \quad (20)$$

3.4 The recursive formulation

After putting all building blocks in place, we present our IAM's recursive formulation with Bayesian learning about the ECS. The social planner chooses capital tomorrow

K_{t+1} and mitigation $\mu_t \in [0, 1]$ to solve

$$V_t(\mathbf{X}_t) = \max_{C_t, K_{t+1}, \mu_t} \left\{ \frac{(C_t/L_t)^{1-1/\psi}}{1-1/\psi} L_t + e^{-\rho} \mathbb{E}_t \left[V_{t+1}(\mathbf{X}_{t+1})^{\frac{1-\gamma}{1-\psi}} \right]^{\frac{1-1/\psi}{1-\gamma}} \right\} \quad (21)$$

$$\text{s.t. } (1 - \Theta(\mu_t)) \Omega_t (T_{AT,t}) K_t^\alpha (A_t L_t)^{1-\alpha} - C_t + (1 - \delta) K_t - K_{t+1} = 0 \quad (22)$$

$$1 - \mu_t \geq 0 \quad (23)$$

$$(1 - b_{12}) M_{AT,t} + b_{21} M_{UO,t} + (1 - \mu_t) \sigma_t K_t^\alpha (A_t L_t)^{1-\alpha} + E_{Land,t} - M_{AT,t+1} = 0 \quad (24)$$

$$b_{12} M_{AT,t} + (1 - b_{21} - b_{23}) M_{UO,t} + b_{32} M_{LO,t} - M_{UO,t+1} = 0 \quad (25)$$

$$b_{23} M_{UO,t} + (1 - b_{32}) M_{LO,t} - M_{LO,t+1} = 0 \quad (26)$$

$$(1 - c_1 c_3 - \varphi_{1C}) T_{AT,t} + c_1 c_3 T_{OC,t} + c_1 \left(F_{2\text{cco2}} \log_2 \left(\frac{M_{AT,t}}{M_{AT}^*} \right) + F_{EX,t} \right) + \varphi_{1C} \tilde{f}_{t+1} T_{AT,t} + \epsilon_{T,t+1} - T_{AT,t+1} = 0 \quad (27)$$

$$c_4 T_{AT,t} + (1 - c_4) T_{OC,t} - T_{OC,t+1} = 0 \quad (28)$$

$$\frac{S_{\epsilon_T} \mu_{f,t} + \varphi_{1C} T_{AT,t} (\varphi_{1C} T_{AT,t} \tilde{f}_{t+1} + \epsilon_{T,t+1}) S_{f,t}}{S_{\epsilon_T} + (\varphi_{1C} T_{AT,t})^2 S_{f,t}} - \mu_{f,t+1} = 0 \quad (29)$$

$$\frac{S_{\epsilon^T} S_{f,t}}{S_{\epsilon^T} + (\varphi_{1C} T_{AT,t})^2 S_{f,t}} - S_{f,t+1} = 0 \quad (30)$$

$$\tilde{f}_{t+1} \sim \mathcal{N}(\mu_{f,t}, S_{f,t}, \underline{f}, \bar{f}), \epsilon_{T,t+1} \sim \mathcal{N}(0, S_{\epsilon^T}) \quad (31)$$

$$\Omega(T_{AT,t}) = 1 - \frac{1}{1 + \left(\frac{1}{20.46} T_{AT,t} \right)^2 + \left(\frac{1}{2 \cdot TP_t} T_{AT,t} \right)^{6.754}} \quad (32)$$

$$TP_{t+1} = TP_t + \epsilon_{TP,t+1}, \quad \epsilon_{TP,t} \sim \mathcal{N}(0, S_{TP}) \quad (33)$$

The Bellman equation is subject to the budget constraint, the law of motion for the capital, the occasionally binding upper bound on mitigation variable μ_t , the three laws of motion of the carbon concentration (in the atmosphere, in the upper ocean and the lower ocean), the two laws of motion of the atmospheric temperature and the ocean temperature, the Bayesian updates of the ECS, and the stochastic law of motion of the damage coefficients.¹⁵

In Eq. (21), the state of the model is given by $\mathbf{X}_t = (K_t, \Gamma_t, \mathbf{S}_t, TP_t, t; \vartheta)$, where $\Gamma_t = (M_{AT,t}, M_{UO,t}, M_{LO,t}, T_{AT,t}, T_{OC,t})$ abbreviates the climate variables, $\mathbf{S}_t = (\mu_{f,t}, S_{f,t})$ the mean and variance of the agent's beliefs, TP_t the climate tipping state (cf. Eq. (7)), and

¹⁵ Notice that the recursive model formulation presented in Eq. (21) is specific for the case of $\psi > 1$. In case of the $\psi < 1$, the value function changes to

$$V_t(\mathbf{X}_t) = \max_{C_t, K_{t+1}, \mu_t} \left\{ -\frac{(C_t/L_t)^{1-1/\psi}}{1-1/\psi} L_t + e^{-\rho} \mathbb{E}_t \left[V_{t+1}(\mathbf{X}_{t+1})^{\frac{1-\gamma}{1-\psi}} \right]^{\frac{1-1/\psi}{1-\gamma}} \right\}.$$

This transformation should be taken into account in all manipulations required to render the model amendable for our solution method (cf. appendix B).

t the state variable that handles time. Finally, ϑ is a collection of the N uncertain parameters $\vartheta = (\vartheta_1, \dots, \vartheta_N)$ in the model which we consider to be pseudo-state variables (Scheidegger and Bilonis, 2019), and over which the planner, in contrast to the other stochastic variables, does not compute expectations. The total dimensionality d of the IAM, therefore, is given by:

$$d = |K_t| + |\Gamma_t| + |S_t| + |TP_t| + |t| + |\vartheta| = 1 + 5 + 2 + 1 + 1 + N = 10 + N. \quad (34)$$

4 Deep uncertainty quantification

In what follows, we introduce our Deep UQ methodology in three steps. First, we outline in Section 4.1 how “Deep Equilibrium Nets” (DEQN) can be adapted to compute global solutions to (non-stationary) stochastic IAMs as a function of its economic and climate states, as well as parameters (Section 4.1.1), and how the resulting optimal policies can subsequently be used to construct cheap-to-evaluate surrogate models for derived quantities of interest, such as the SCC, by using the GP-based surrogate model (Section 4.1.2). Section 4.2 introduces the global sensitivity measures we intend to study in the context of our IAM. Section 4.3 finally presents the Deep UQ algorithm formally. Additional details on the implementation are provided in Appendix B.

4.1 Deep surrogate models for IAMs

4.1.1 Deep equilibrium nets

In this section, we briefly summarize the general idea of DEQNs, thereby adopting the notation of Azinovic et al. (2022), Bretscher et al. (2022). The DEQN algorithm is a simulation-based solution method using deep neural networks to compute an approximation of the *optimal policy function* $\mathbf{p} : X \rightarrow Y \subset \mathbb{R}^M$ to a dynamic model under the assumption that the underlying economy can be characterized via discrete-time first-order equilibrium conditions, that is,

$$\mathbf{G}(\mathbf{x}, \mathbf{p}) = \mathbf{0}, \quad \forall \mathbf{x} \in X \subset \mathbb{R}^d. \quad (35)$$

Intuitively, DEQNs work as follows: An unknown policy function is approximated with a neural network, that is, $\mathbf{p}(\mathbf{x}) \approx \mathcal{N}(\mathbf{x})$ with trainable parameters ν , which are ex-ante unknown and that have to be determined based on some suitable loss function measuring the quality of a given approximation at a given state of the economy.

Although there are several different types of deep neural networks, in this paper, we use the so-called densely-connected feedforward neural networks (FNN).¹⁶ Following the literature, we define an L -layer FNN as a function $\mathcal{N}^L(\mathbf{x}) : \mathbb{R}^{d_{\text{input}}} \rightarrow \mathbb{R}^{d_{\text{output}}}$ and say that there are $L - 1$ hidden layers such that the ℓ -th layer has N_ℓ neurons. In our concrete

¹⁶ Neural networks are universal function approximators (Hornik et al., 1989) that can resolve highly non-linear features, and that can handle a large amount of high-dimensional input data. See, for example, Goodfellow et al. (2016) for a general introduction to deep learning, and Scheidegger et al. (2023) for a specific introduction in the context of economics.




Figure 1: The figure above depicts a stylized FNN with an input \mathbf{x} , which is a 12-dimensional input vector. It consists of two hidden layers, each containing 16 neurons, and an output $\mathbf{p}(\mathbf{x})$, which is an 11-dimensional vector. This configuration represents a stylized architecture for an IAM that incorporates 10 economic and climate state variables, two additional parameters treated as pseudo-states, and is designed to generate eleven policies.

case, $K = N_0 = d_{\text{in}}$ and $M = N_L = d_{\text{output}}$.¹⁷ Furthermore, for each $1 \leq \ell \leq L$, we define a so-called weight matrix $\mathbf{W}^\ell \in \mathbb{R}^{N_\ell \times N_{\ell-1}}$ and bias vector $\mathbf{b}^\ell \in \mathbb{R}^{N_\ell}$. Then, letting $A^\ell(\mathbf{x}) = \mathbf{W}^\ell \mathbf{x} + \mathbf{b}^\ell$ be the affine transformation in the ℓ -th layer, for some non-linear activation function $\sigma(\cdot)$ such as relu, swish, or selu, an FNN is given by

$$\mathbf{p}(\mathbf{x}) \approx \mathcal{N}(\mathbf{x}) = \mathcal{N}^L(\mathbf{x}) = A^L \circ \sigma_{L-1} \circ A^{L-1} \circ \dots \circ \sigma_1 \circ A^1(\mathbf{x}). \quad (36)$$

In Figure 1, we illustrate a simple FNN with two hidden layers. The selection of hyper-parameters $\{L, \{N_\ell\}_{\ell=1}^L, \{\sigma_\ell(\cdot)\}_{\ell=1}^L\}$ is known as the architecture selection. Approaches to determine these hyper-parameters include using prior experience, manual, random, or grid search, as well as more complex methods such as Bayesian optimization (see, e.g., [Bergstra et al. \(2011\)](#)).

The DEQN algorithm to determine $\mathbf{p}(\mathbf{x})$ is started by randomly initializing the v 's ([Glorot and Bengio, 2010](#)), that is, an arbitrary guess for the ex-ante unknown approximate policy function. Next, we simulate a sequence of $N_{\text{path length}}$ states. Starting from

¹⁷ Recall that in our model, $d_{\text{in}} = d = 10 + N$ input dimensions (cf. Eq. (34)). Furthermore, $d_{\text{output}} = 3 + 7 + 1 = 11$, that is, two choice variables (consumption, capital tomorrow, mitigation, seven Lagrange multipliers, and the value function (see Section 3.4 and appendix B.1 for more details).

some given state \mathbf{x}_t , the next state \mathbf{x}_{t+1} is the result of the policies encoded by the neural network, $\mathcal{N}(\mathbf{x})$, and remaining model-implied dynamics.

If we knew the (approximate) policy function satisfying the equilibrium conditions, Eq. (35) would hold along a simulated path. However, since the neural network is initialized with random coefficients, $\mathbf{G}(\mathbf{x}_t, \mathcal{N}(\mathbf{x}_t)) \neq 0$ along the simulated path of length $N_{\text{path length}}$. This fact is now leveraged to improve the quality of the guessed policy function. Specifically, DEQNs use a loss function as the error in the equilibrium conditions, that is,

$$\ell_{\nu} := \frac{1}{N_{\text{path length}}} \sum_{\mathbf{x}_t \text{ on sim. path}} \sum_{m=1}^{N_{\text{eq}}} (\mathbf{G}_m(\mathbf{x}_t, \mathcal{N}(\mathbf{x}_t)))^2, \quad (37)$$

where $\mathbf{G}_m(\mathbf{x}_t, \mathcal{N}(\mathbf{x}_t))$ represent all the N_{eq} first-order equilibrium conditions of a given model, that is, $\mathbf{G}(\mathbf{x}_t, \mathcal{N}(\mathbf{x}_t)) = \sum_{m=1}^{N_{\text{eq}}} (\mathbf{G}_m(\mathbf{x}_t, \mathcal{N}(\mathbf{x}_t)))$. Eq. (37) can now be used to update the weights of the network with any variant of (stochastic) gradient descent,¹⁸ namely,

$$v'_k = v_k - \alpha^{\text{learn}} \frac{\partial \ell(\nu)}{\partial v_k}, \quad (38)$$

where v'_k represents the updated k -th weight of the neural network, and where $\alpha^{\text{learn}} \in \mathbb{R}$ denotes the so-called learning rate. The updated neural network-based representation of the policy is subsequently used to simulate a sequence of length $N_{\text{path length}}$ steps, along which the loss function is recorded, and the latter is again used to update the network parameters. This iterative procedure is pursued until $\ell_{\nu} < \epsilon \in \mathbb{R}$, that is, an approximate equilibrium policy, has been found.

In summary, the DEQN algorithm consists of four building blocks: i) deep neural networks for approximating the equilibrium policies; ii) a suitable loss function measuring the quality of a given approximation at a given state of the economy; iii) an updating mechanism to improve the quality of the approximation; and iv) a sampling method for choosing states for updating and evaluating of the approximation quality. In Appendix B.1, the step-by-step procedure for mapping stochastic and non-stationary climate economic models in general, and our model in particular (cf. Section 3.4), onto the neural network-based DEQN solution framework is provided.

4.1.2 Gaussian process surrogate models for global sensitivity analysis

We now briefly introduce the Gaussian process regression (GPR), which is a probabilistic approach for modeling a regression problem¹⁹. We will use this method below to construct surrogates for approximating and interpolating QoIs, such as the SCC, as a function of the model input parameters, so as to enable the swift computation of global sensitivity measures such as the Sobol' indices, the univariate effects, and the Shapley values (cf. Section 4.2 below).

GPR assumes that the underlying function is a sample from a GP. The latter is

¹⁸ In our practical applications, we use "Adam" (Kingma and Ba, 2014).

¹⁹ See, e.g., Rasmussen and Williams (2005) for a textbook treatment.

defined by a mean function, $m(x)$, and a covariance function, $k(x, x')$. For a regression problem, we are given a set of n input-output pairs, $\mathcal{D} = \{x_i, y_i\}_{i=1}^n = [X, y]$, where $x_i \in \mathbb{R}^d$, $y_i \in \mathbb{R}$, $X \in \mathbb{R}^{n \times d}$ and $y \in \mathbb{R}^{n \times 1}$, and to which the literature also refers to as “training dataset”, or “experimental design”.²⁰ The goal is to “learn”, that is, approximate a function $f(x)$ that maps inputs to outputs. In GPR, one often assumes that the output y is generated by a noisy evaluation of the function $f(x)$, that is, $y = f(x) + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$, as it renders the computations numerically more stable (see, e.g., [Rasmussen and Williams \(2005\)](#), [Murphy \(2012\)](#)). In order to approximate functions efficiently with GPs in general, we need to find a set of points x_i that covers the relevant state space in such a way that a GP can approximate the desired QoI, such as the SCC as a function of the model parameters, sufficiently well with a minimal amount of training data. The task at hand, thus, is to answer the question of what data one should gather to learn about the function(s) of interest as quickly as possible, especially in our case where training data is expensive to acquire, as we need to simulate, for instance, the SCC in the year 2100. Various methods have been proposed to achieve this, such as (nested) Latin hypercube sampling (see, e.g., [Harenberg et al. \(2019\)](#)), or Bayesian active learning (see, e.g., [Renner and Scheidegger \(2018\)](#), and references therein). For more details on how we achieve this in our practical applications, see Appendix C.2.

The prior distribution over functions is given by a GP, that is, $f(x) \sim \mathcal{GP}(m(x), k(x, x'))$. In practical applications, the mean function, $m(x)$, is often set to a constant value, for example, zero. The covariance function, $k(x, x')$, defines the correlation between two arbitrary inputs x and x' . The exact choice of the kernel within an application depends on how the modeler encodes prior knowledge about the function(s) to be approximated. In our work below, we use the Matérn 5/2 kernel, which is given by:

$$k(x, x') = \sigma_f^2 \left(1 + \frac{\sqrt{5}r}{\rho} + \frac{5r^2}{3\rho^2} \right) \exp\left(-\frac{\sqrt{5}r}{\rho}\right), \quad (39)$$

where σ_f^2 denotes the variance parameter, and ρ represents the length scale parameter. The distance r between x and x' can be calculated using a suitable metric, such as the Euclidean distance.²¹

Given the input-output pairs in \mathcal{D} , we can compute the posterior distribution over functions using Bayes’ rule, that is,

$$p(f_* | x_*, \mathcal{D}) = \mathcal{N}\left(\mu_*, \sigma_*^2\right), \quad (40)$$

where x_* is an arbitrary point from the computational domain (also referred to as “test input”), f_* is the corresponding output, and μ_* and σ_*^2 are the mean and variance of

²⁰ Below, the input vector x_i will be given by the vector of parameters of dimensionality N (cf. Eq. (34)), whereas corresponding observations y_i will represent the QoIs.

²¹ Note that the hyperparameters of the covariance function are typically estimated by maximizing the likelihood ([Rasmussen and Williams, 2005](#)). For more details on kernels, see, for example, [Murphy \(2022, Ch. 18.2\)](#).

the predictive distribution. They can be computed as:

$$\mu_*(x_*) = k(x_*, X) [K(X, X) + \sigma_\epsilon^2 I]^{-1} y, \quad (41)$$

$$\sigma_*^2(x_*) = k(x_*, x_*) - k(x_*, X) [K(X, X) + \sigma_\epsilon^2 I]^{-1} k(X, x_*), \quad (42)$$

where $K(X, X')$ is the matrix of pairwise covariances between inputs in X and X' . Furthermore, I is an identity matrix and σ_ϵ is the assumed noise level of observations, that is, the variance of ϵ . Thus, the interpolation value of a function at a location x_* is given by $f(x_*) = \mu_*$. For more details on GPs, see, e.g., [Renner and Scheidegger \(2018\)](#), and references therein.

Below, we need cheap-to-evaluate GP surrogates for one particular QoI as a function of the model input parameters, the SCC, as it is relevant for the climate policy, and for which we intend to perform GSA at any given point in time. Recall that the SCC is defined as the marginal cost of atmospheric carbon in terms of the numeraire good. Following the literature (see, e.g., [Traeger \(2014\)](#)), we define SCC as the marginal rate of substitution between the atmospheric carbon concentration of a 1000 Gt of carbon where $c2co2$ represents the carbon to CO_2 transformation coefficient²² and the normalized capital stock, that is,

$$SCC_t = \frac{-(\partial V_t / \partial M_{AT,t}) / c2co2}{\partial V_t / \partial K_t} = \frac{-(\partial v_t / \partial M_{AT,t}) / c2co2}{\partial v_t / \partial k_t} A_t L_t. \quad (43)$$

To enable a comparison to related studies, we focus below on the SCC in the year 2100, a date which is commonly used in the literature (see, e.g., [Cai and Lontzek \(2019\)](#)), that is,

$$y_{SCC} = \mathbb{E} [SCC_{2100}]. \quad (44)$$

To construct GP surrogates of the QoI mentioned above as a function of the model input parameters ϑ , we proceed as follows:²³ Suppose that we observe n model input-output pairs, we obtain a training set (which in the machine learning literature is often termed experimental design), $\mathcal{D} = \{\vartheta_i, y_i\}_{i=1}^n = [\Theta, y]$ consisting of n sample points $\vartheta_i \in \mathcal{B} \subset \mathbf{R}^N$ from a joint (uniform) distribution of the uncertain model parameters, whereas $y_i \in \mathbb{R}$ corresponds to the QoI at a given combination of parameters and can be generated by interpolating on the DEQN surrogate of the pre-computed IAM solution (cf. Section 3.4). Next, we fit a GP model to the training data set \mathcal{D} , and denote it as $\mathcal{M}_{GP|\Theta, y}$. We use the predictive mean, $\mu_*(\vartheta)$, of the respective GP model as the interpolation value for the respective QoI, that is, y_{SCC} .²⁴ Evaluating $\mu_*(\cdot)$ numerically

²² The stock of carbon in the carbon cycle is measured in 1000 Gt of carbon. However, the backstop cost, as in [Nordhaus \(2017\)](#) is measured in 1000 Gt of CO_2 emissions. Thus, to obtain a correct value for SCC, we need to transform a value measured in carbon into CO_2 by the transformation coefficient $c2co2 = 3.666$.

²³ In Section 4.1.2, we generally introduce the GPR that is trained on a set of n input-output pairs, which is $\mathcal{D} = \{x_i, y_i\}_{i=1}^n = [X, y]$. Hereafter we focus more on our specific case where the model input parameters are denoted by ϑ_i , not x_i and define the n input-output pairs $\mathcal{D} = \{\vartheta_i, y_i\}_{i=1}^n = [\Theta, y]$.

²⁴ Note that the GP surrogate is redundant when a QoI is not a derived quantity from the optimal policies such as μ_t . In contrast, where a QoI relies on simulations such as the SCC, constructing a

is extremely fast, as the algorithmic complexity is linearly in size n of the training data set (cf. Eq. (41)). Thus, GSA of very rich specified models (cf. Section 4.2) now becomes numerically tractable. More details on how we assess the accuracy of GP surrogate models are provided in Appendix C.1, whereas in appendix C.2, we describe how we efficiently choose training data.

4.2 Global sensitivity analysis in a nutshell

In this section, we formally introduce surrogate-based global sensitivity analysis (GSA).

4.2.1 Some definitions

We begin our discussion of GSA by introducing essential notation, thereby following [Saltelli et al. \(2007\)](#), [Sudret \(2008\)](#) and [Harenberg et al. \(2019\)](#). We define a (true) mathematical model $\mathcal{M}(\cdot)$ that maps

$$\vartheta \in \mathcal{D}_\vartheta \subset \mathbb{R}^N \rightarrow y = \mathcal{M}(\vartheta) \in \mathbb{R}, \quad (45)$$

where y is called a quantity of interest (QoI)²⁵ that is, a random endogenous outcome of the computational model $\mathcal{M}(\cdot)$, where $\vartheta = (\vartheta_1, \dots, \vartheta_N)$ is the random vector of N input parameters, where each parameter is characterized by a probability density function (PDF) f_{ϑ_i} , $i = 1, \dots, N$, and where the joint density f_ϑ is defined over a probabilistic space (see, e.g., [Jaynes \(1957, 1982\)](#)). The corresponding ex-ante unknown distribution of the endogenous output y is inferred by evaluating the model on a large sample of parameter values drawn from the specified distribution f_ϑ ; a technique that is termed *uncertainty propagation* (see, e.g., [Sudret \(2008\)](#) and references therein). Based on this distribution, we next look at three particular types of metrics that are typically studied in the UQ literature, namely, the Sobol' indices (Section 4.2.2), the univariate effects (Section 4.2.3), and the Shapley values (Section 4.2.4).

4.2.2 Sobol' indices

Variance-based GSA methods discriminate among all the model's parameters according to their contribution to the variance of its output. In our concrete case, we intend to study which (combinations of) parameters drive the QoI's variance by applying Sobol's decomposition ([Sobol, 2001](#)). Following a common practice in the UQ literature, for instance, see [Saltelli et al. \(2007\)](#), we start from the Sobol's variance decomposition to investigate possible interactions among model inputs on the model output y , that is,²⁶

GP-based surrogate can reduce the computational burden substantially since apart from evaluating the exact formula of the SCC at n points to train a GP, everywhere else in the parameter space, the respective quantity can be retrieved by evaluating the surrogate model.

²⁵ Recall that in our case, we consider the SCC as a function of the model parameters at a given point in time (cf. Eq. (43)).

²⁶ Sobol's decomposition is closely related to a function approximation technique called "high-dimension model representation" (see, e.g., [Ma and Zabarabaz \(2010\)](#), [Eftekhari et al. \(2017\)](#), [Eftekhari and Scheidegger \(2022\)](#)).

$$\text{Var}[y] = \sum_{i=1}^N V_i + \sum_{1 \leq i < j \leq N} V_{i,j} + \cdots + V_{1,2,\dots,N}, \quad (46)$$

where V_u denotes the partial variances for any subset of parameter indices $u \subset \{1, 2, \dots, N\}$, and in particular in our case,

$$V_i = \text{Var}[\mathbb{E}[\mathcal{M}(\boldsymbol{\vartheta}) \mid \vartheta_i]], \quad (47)$$

and

$$V_{i,j} = \text{Var}[\mathbb{E}[\mathcal{M}(\boldsymbol{\vartheta}) \mid \vartheta_i, \vartheta_j]] - \text{Var}[\mathbb{E}[\mathcal{M}(\boldsymbol{\vartheta}) \mid \vartheta_i]] - \text{Var}[\mathbb{E}[\mathcal{M}(\boldsymbol{\vartheta}) \mid \vartheta_j]], \quad (48)$$

and so forth.

A natural extension to the relative shares of partial variances in the total variance leads to the well-known sensitivity measures, that is, Sobol' indices (Sobol, 2001). For any subset of parameters' indices u , the Sobol' index is defined as

$$S_u = \frac{\text{Var}_{\vartheta_u}[\mathbb{E}_{\boldsymbol{\vartheta} \setminus \vartheta_u}[\mathcal{M}(\boldsymbol{\vartheta}) \mid \vartheta_u]]}{\text{Var}_{\boldsymbol{\vartheta}}[Y]}. \quad (49)$$

The Sobol' indices given by Eq. (49) are indicators used in the variance-based sensitivity analysis (Sobol, 2001). In practical terms, Sobol' indices quantify which uncertain parameters and non-linear interactions among uncertain parameters primarily contribute to the variances of model outcomes. In other words, one can *screen* uncertain parameters to decrease overall model variance by investigating the Sobol' indices. In our work below, we particularly look at the first-order Sobol' index when we set $u = \{i\}$ in Eq. (49), that is,

$$S_i = \frac{\text{Var}_{\vartheta_i}[\mathbb{E}_{\boldsymbol{\vartheta} \setminus \vartheta_i}[\mathcal{M}(\boldsymbol{\vartheta}) \mid \vartheta_i]]}{\text{Var}_{\boldsymbol{\vartheta}}[Y]}. \quad (50)$$

We follow, among others, Oakley and O'Hagan (2004), Marrel et al. (2009), and use the predictive mean of the fitted Gaussian process (GP) surrogate model $\mathcal{M}_{\text{GP}|\Theta_s, Y_s}$ instead of the "true" $\mathcal{M}(\cdot)$ to enable a swift computation of GSA metrics.²⁷ To compute the first-order Sobol' indices defined in Eq. (50), we start by generating test data inputs $\Theta = \{\boldsymbol{\vartheta}_i\}_{i=1}^{N_\omega}$ on which we simulate the GP model $\mathcal{M}_{\text{GP}|\Theta, Y}$ for N_ω times to obtain the model predictions. The sample space Ω represents the possible model outcomes ω from the posterior. Thus, using the predictive mean of the GP, the first-order Sobol' indices for the parameter i are given by (Marrel et al., 2009)

$$S_i = \frac{\text{Var}_{\boldsymbol{\vartheta}_i} \mathbb{E}_{\Theta \setminus \boldsymbol{\vartheta}_i}[\mathbb{E}_{\Omega}[\mathcal{M}_{\text{GP}|\Theta, Y}(\Theta) \mid \boldsymbol{\vartheta}_i]]}{\text{Var}_{\Theta} \mathcal{M}_{\text{GP}|\Theta, Y}(\Theta)}. \quad (51)$$

²⁷ Apart from GPs, polynomial chaos expansion is also commonly used in the literature to construct surrogate models for the swift computation of Sobol' indices (see, e.g., Sudret (2008)).

4.2.3 Univariate effects

From the aforementioned variance decomposition, we do not obtain any information as to which direction a QoI moves when the uncertain parameters deviate from their benchmark values. To this end, GSA is also concerned with computing univariate effects (Younes et al., 2013). Formally, a univariate effect represents the model’s conditional expectation of a QoI when we fix a single model parameter ϑ_i at an arbitrary point ϑ'_i , where the expectation value is taken over all other parameters, that is,

$$\mathcal{U}_i(\vartheta'_i) = \mathbb{E}_{\vartheta \setminus \vartheta_i} [\mathcal{M}(\vartheta) \mid \vartheta_i = \vartheta'_i]. \quad (52)$$

As before, we replace the true “true” $\mathcal{M}(\cdot)$ in Eq. (52) again by the GP surrogate model $\mathcal{M}_{\text{GP}|\vartheta,y}$, and evaluate the model outputs by using the posterior mean of the GP.

4.2.4 Shapley value

An essential task in UQ is to attribute the uncertainty of the overall outcome of the model to various input parameters. In the field of GSA, the Sobol’ indices have been widely adopted for variance decomposition (Saltelli et al., 2007, Harenberg et al., 2019). On the other hand, the economic literature, especially in cooperative game theory, has studied a similar problem and developed a solution concept known as the Shapley value (see, e.g., Shapley (1953), Winter (2002)). The key idea behind the Shapley value is to determine a fair way to attribute the value created by a team effort to each team member. Based on this insight, Owen (2014) introduced a global sensitivity measure based on the Shapley value, which is a compelling choice for identifying how much model variance can be attributed to the uncertainty in input parameters.

Following Owen (2014), Song et al. (2016) and our previously introduced notation, we define the Shapley value for the i -th uncertain parameter as:

$$Sh_i = \sum_{\mathbf{u} \subseteq \mathcal{K} \setminus i} \frac{(N - |\mathbf{u}| - 1)! |\mathbf{u}|!}{N!} (\text{Var} [\mathcal{M}(\vartheta_{\mathbf{u} \cup \{i\}})] - \text{Var} [\mathcal{M}(\vartheta_{\mathbf{u}})]), \quad (53)$$

where $\mathcal{K} = \{1, 2, \dots, N\}$, and $N = |\mathcal{K}|$ is the size of the whole set and $|\mathbf{u}|$ is the size of subset \mathbf{u} . We also define $\text{Var} [\mathcal{M}(\vartheta_{\emptyset})] = 0$. Eq. (53) measures the incremental cost when including parameter i in set \mathbf{u} averaged over all sets $\mathbf{u} \subseteq \mathcal{K} \setminus i$. The variance in Eq. (53) is

$$\text{Var} [\mathcal{M}(\vartheta_{\mathbf{u}})] = \mathbb{E}_{\mathcal{K} \setminus \mathbf{u}} [\text{Var}_{\mathbf{u}} [\mathcal{M}(\vartheta) \mid \vartheta_{\mathcal{K} \setminus \mathbf{u}}]] \quad (54)$$

with which one can measure a similar effect as that of the first-order Sobol’ indices.²⁸

²⁸ Alternatively, Owen (2014) proposes to choose

$$\text{Var} [\mathcal{M}(\vartheta_{\mathbf{u}})] = \text{Var}_{\vartheta_{\mathbf{u}}} [\mathbb{E}_{\mathcal{K} \setminus \mathbf{u}} [\mathcal{M}(\vartheta) \mid \vartheta_{\mathbf{u}}]]. \quad (55)$$

Song et al. (2016) prove that the Shapley values using either the variance definition in Eq. (54) or Eq. (55) are equivalent. Furthermore, when numerically estimating the Shapley value, Sun et al. (2011) pointed out that the Monte Carlo estimation of the inner expectation operator in Eq. (55) requires a sufficiently large number of model evaluations and could be highly biased. On the other hand, the estimation in

When the number of uncertain parameters N , as in our case, is reasonably “small”, one typically uses the “exact” method by [Song et al. \(2016\)](#). The latter involves traversing $N!$ permuted sets. To enhance the efficiency of the Monte Carlo sampling required for precise evaluation of the expectation and variance operators, we again employ the GP-based surrogate model, thereby replacing the original model, that is, we utilize the predictive mean of the posterior for this purpose ([Iooss and Prieur, 2019](#)).

4.3 The deep uncertainty quantification algorithm

After having discussed the ingredients to our Deep UQ methodology in Section 4.1 and Section 4.2, we now present code listing 1 the complete algorithm. It consists of two fundamental building blocks. First, the DEQN is used to solve the IAM as a function of all endogenous, exogenous, and pseudo-state variables in a single solution step. The resulting neural network-based surrogate model then serves as an input to construct GP surrogates for QoI metrics, such as the SCC. Subsequently, UQ can be performed on these GP surrogates to assess and analyze the uncertainty in the model predictions.

Eq. (54) is unbiased for all sample sizes.

Algorithm 1 Deep UQ algorithm

Part 1: Deep equilibrium nets

Approximate policy functions $p(x_t)$ using DEQN parameterized by the coefficients ν for the state x_t .

Input:

Neural network architecture and hyper-parameters, simulation length T (length of an episode), α^{learn} (learning rate), x_0 (initial states). Upper and lower bound of the pseudo states $\vartheta \in [\underline{\vartheta}, \bar{\vartheta}] \subset \mathbb{R}^N$.

Output:

The neural network's trained parameters ν^* , such that the policy function $p(x_t) \approx \mathcal{N}(x_t)$ minimizes the loss function, i.e., the system of the first-order equilibrium conditions.

```

i = 1
x0i ← x0                                     ▶ Starting point of simulation
while  $\ell_\nu > \epsilon$  do
     $\mathcal{D}_{\text{train}}^i \leftarrow \{x_0^i, x_1^i, x_2^i, \dots, x_T^i\}$    ▶ Generate training data
    Evaluate the loss function along a simulated path:

    
$$\ell_\nu := \frac{1}{N_{\text{path length}}} \sum_{x_t \text{ on sim. path}} \sum_{m=1}^{N_{\text{eq}}} \left( G_m(x_t^i, \mathcal{N}(x_t^i)) \right)^2$$


    for  $k \in \{1, 2, \dots, \text{len}(\nu)\}$  do               ▶ Perform stochastic gradient descent

        
$$\nu_k^{i+1} = \nu_k^i - \alpha^{\text{learn}} \frac{\partial \ell(\nu^i)}{\partial \nu_k^i}$$


    end for
    i ← i + 1                                       ▶ Go to the next episode i + 1
     $\nu^i \leftarrow \nu^{i+1}$                              ▶ Update the neural network that parameterizes the policies
    x0i ← xTi-1                                   ▶ Starting point for the simulation in the next episode
end while
return  $\nu^* \leftarrow \nu^i$                              ▶ Optimal policy at convergence:  $p^*(x_t) \approx \mathcal{N}(x_t)$ 

```

Part 2: Uncertainty quantification based on GP surrogate models

Uncertainty quantification using the predictive mean of the GP surrogate model $\mathcal{M}_{\text{GP}|\vartheta, y}$.

Input:

$p^*(x_t)$ (optimal policy functions), n (size of the training set for the GPs), and y (QoI, e.g., the SCC).

Output:

S_i (the first-order Sobol' indices), U_i (the univariate effects) and Sh_i (the Shapley values) for the uncertain parameter ϑ_i .

```

while  $\epsilon_{\text{GP}|\vartheta, y}^{\text{LOO}} \geq \epsilon_{\text{GP}}^{\text{LOO}}$  do ▶ Criterion to ensure a high-quality GP surrogate model (cf. appendix C.1)
    for  $j \in \{1, \dots, n\}$  do                             ▶ Generate a training set  $\mathcal{D}$  for the GPs
         $y_j = \mathcal{M}(\vartheta_j)$  ▶ Compute the QoI for a given combination of parameters  $\vartheta_j$  by using  $p^*(x_t)$ 
    end for
    Fit the GP model  $\mathcal{M}_{\text{GP}|\vartheta, y}$  to  $\mathcal{D}$ 
    Compute the leave-one-out error  $\epsilon_{\text{GP}|\vartheta_s, y_s}^{\text{LOO}}$ 
    n ← n + 1                                             ▶ Increase the size of the training set if necessary
end while
return  $\mathcal{M}_{\text{GP}|\vartheta, y}$  ▶ Accurate GP surrogate model; used to replace true model  $\mathcal{M}(\cdot)$  in UQ tasks
for  $i \in \{1, \dots, N\}$  do                               ▶ UQ for each uncertain parameter  $\vartheta_i$ 
    Compute the Sobol' indices  $S_i$  (cf. Eq. (51)).
    Compute the univariate effects  $\mathcal{U}_i$  (cf. Eq. (52)).
    Compute the Shapley values  $Sh_i$  following (cf. Eq. (53)).
    return  $S_i, \mathcal{U}_i$  and  $Sh_i$ 
end for

```

5 Results

In this section, we apply our Deep UQ method (cf. Section 4) to study the stochastic IAM we have posited in Section 3. It is important to note that the results presented here are preliminary. The model solutions we present below are produced using the damage function detailed in Eq. (5), and do not take into account tipping points. After describing in Section 5.1 the representative set of parameters over which we perform UQ, we proceed in two steps. First, we briefly present in Section 5.2 the optimal policy results for two variations of our IAM: one model that contains solely a temperature shock (but no learning about the ECS), and our benchmark model, which incorporates Bayesian learning over the ECS. In addition, we discuss the UQ results for the two variations of our IAM. Second, we study in Section 5.3 how fat upper-tailed uncertainty over the temperature change from a doubling of GHG could be diminished through Bayesian learning.

5.1 Parameters for the uncertainty quantification analysis

To strike a balance across the uncertainty embedded in the economic and climate block in our IAM, we perform UQ over six subjectively selected uncertain parameters $\vartheta \in \mathbb{R}^6$. Their ranges are summarized in Table 1: three of them reflect (parametric) uncertainty in the economic part of the IAM, whereas the other three are from the climate module. From the economic module, we include the pure rate of time preferences ρ , the Arrow-Pratt relative risk aversion γ and the intertemporal elasticity of substitution ψ (IES) that characterizes the Epstein-Zin preferences. As it is discussed in [Ju and Miao \(2012\)](#), [Basu and Bundick \(2017\)](#), [De Groot et al. \(2018\)](#), computations in the neighborhood of IES being unity are numerically unstable. Thus, to avoid this numerical issue, we split the interval of interest for the IES into two subintervals, thereby avoiding unity.

From the climate module of our IAM, we include the initial values of the prior mean $\mu_{f,0}$ and variance $S_{f,0}$ of the Bayesian learning process over ECS in our UQ exercise. Furthermore, we also include the damage parameter π_2 that is a part of the quadratic damage function (cf. Eq. (5)).

As there is no a priori knowledge available on the distribution of each uncertain parameter ϑ_i , in line with [Harenberg et al. \(2019\)](#), we assume a uniform distribution with a lower $\underline{\vartheta}_i$ and upper bound $\bar{\vartheta}_i$ around benchmark values ϑ_i^0 typically used in the literature. The subjectively, yet plausible ranges for the statistically independent parameters $\vartheta_i \in [\underline{\vartheta}_i, \bar{\vartheta}_i]$ are reported in Table 1.

5.2 Uncertainty quantification results

This section discusses the optimal policy results and the UQ for the IAM introduced in Section 3. The results of this model with Bayesian learning over the ECS are contrasted with computations of a simplified model that features a random ECS and temperature shock but assumes the agent does not learn from it. With this comparison, we intend to clarify the impact of learning on the optimal policies and its implications on UQ. One should think about the model with a random ECS as a case where the social

Table 1: Parameter ranges for the UQ analysis.

Parameter	ϑ_i^0	$\underline{\vartheta}_i$	$\overline{\vartheta}_i$	Source of the parameter value
ρ	0.015	0.01	0.02	Stern (2008)
γ	10.0	5.0	10.	Jensen and Traeger (2014) and Cai and Lontzek (2019)
$\psi < 1$	0.5	0.5	0.95	Jensen and Traeger (2014) and Cai and Lontzek (2019)
$\psi > 1$	1.5	1.05	2.0	Jensen and Traeger (2014) and Cai and Lontzek (2019)
π_2	0.00236	0.002	0.008	Nordhaus (2017) and Weitzman (2012)
$\mu_{f,0}$	0.65	0.45	0.73	Roe and Baker (2007) and Folini et al. (2023)
$S_{f,0}$	0.0169	0.01	0.04	Roe and Baker (2007)




Figure 2: Emissions as a function of years, ranging from the year 2015 - 2100, for the IAM without learning (left panel), and with learning (right panel), and generated with 1,000 sample paths.

planner has prior knowledge about the mean value of the uncertain climate feedback parameter that enters the ECS (cf. Eq. (14)) and knows the distribution. The social planner considers the stochasticity of the temperature in decision-making, but there is no learning about it in the sense that there is no Bayesian update on the prior values of the distribution of the climate feedback parameter.

In the remainder of this article, we will, for the sake of brevity, call the model setup with a stochastic ECS and temperature shock as the “no-learning case”, and our benchmark model with Bayesian learning as the “learning case”. Furthermore, unless indicated otherwise, we present the results for the benchmark values of the uncertain parameters, denoted as ϑ_i^0 in Table 1, and $\psi > 1$, mostly fixed at $\psi = 1.5$.

Figure 2 depicts the industrial emission paths for the no-learning (left panel) and learning cases (right panel). We can see that in the presence of learning, the amount of emissions significantly increases. The reason for this effect is that the social planner learns the ECS more efficiently if there is more carbon in the atmosphere. Thus, an increased amount of carbon in the atmosphere raises the signal-to-noise ratio, and in turn, the agent can learn quickly. However, the increase of carbon in the atmosphere also raises the temperature and, in turn, leads to higher damages. Thus, in the optimal case, the agent must trade off a necessary increase in emissions to learn quickly, as well

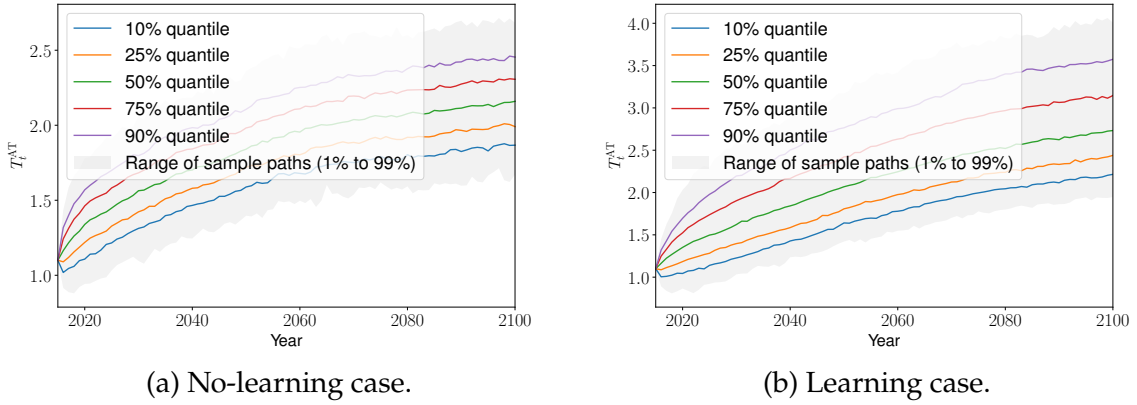


Figure 3: Atmospheric temperature increase for the years 2015 - 2100 for the IAM without learning (left panel), and with learning (right panel), generated with 1,000 sample paths.

as the damages that arise from this increase. Figure 2 shows this necessary increase in emissions for the learning to happen compared to the no learning case.

Figure 3 demonstrates that the increase in emissions necessary to facilitate learning is also reflected in the atmospheric temperature levels. The ability to learn about the ECS comes at the cost of roughly one degree more warming in the year 2100 on average.

From the findings discussed so far, it follows that the SCC shown in Figure 4 is lower for the learning case than the no-learning case. Furthermore, we can also observe that there is almost no variance in the SCC in the no-learning case, but a significant variance in the SCC in the learning case. The variance of the SCC in the learning case can be explained as follows: The noise in the learning process stems from the temperature shock. If the agent is unlucky and gets hit by a sequence of high-temperature, that is, weather shocks (see, for example, the 90th quantile in the Figure 3b which translates into the respective 90th quantile in the Figure 4b), the noise is very high, and consequently, there is a need for more emissions to learn about the ECS value. If the weather shocks are low (see, for example, the 10th quantile in Figure 3b and Figure 4b), there is little noise, and the agent learns quickly enough without excessive emissions necessary for increasing signal-to-noise ratio in the learning process. In contrast, when the agent cannot learn about the ECS value, she mitigates in any case to insure against the uncertain temperature increase. Overall, the results from comparing the two types of models confirm the findings of the previous literature (see, e.g., Kelly and Tan (2015), Hwang et al. (2017)).

Figure 5 shows the sensitivity of the SCC value with respect to three different levels of the IES. The three panels confirm the previous results by Jensen and Traeger (2014) as well as findings by Cai and Lontzek (2019), who report a rise in the SCC as a response to an increase in the IES.

Next, due to the computational advantages of our Deep UQ methodology, we can explore the sensitivity of the SCC (or any other QoI) with respect to any combination of the uncertain parameters with a thorough UQ analysis (cf. Section 4) at negligible

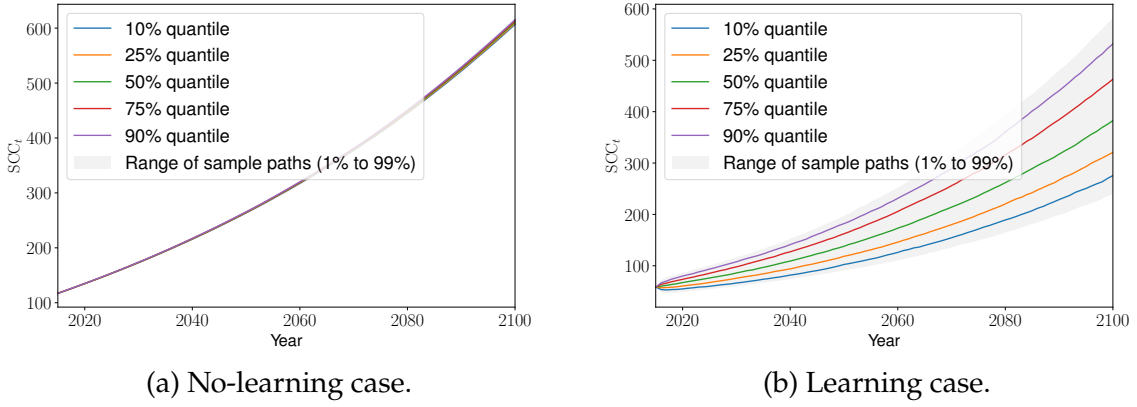


Figure 4: SCC for the years 2015 - 2100 for the IAM without learning (left panel), and with learning (right panel), generated with 1,000 sample paths.




Figure 5: SCC as a function of years ranging from the year 2015 - 2100 for the IAM with Bayesian learning and different levels of the IES, generated with 1,000 sample paths.

computational costs. In Figure 6, we provide the Shapley values (in blue) and Sobol' indices (in green) for the SCC in the years 2020, 2050, and 2100 for the model with no learning. Recall that both of those UQ measures consider the contribution of the variance of the uncertain parameters to the variance of the QoI – in our case – the SCC. For the IAM with no learning, the predominant factor influencing the variability in the SCC is the uncertainty associated with the ECS, represented by the μ_f . The second most significant contributor to this UQ measure is the parameter quantifying climate-related damages, which is labeled as π_2 .

In Figure 7, we present the UQ results for the IAM with learning. We can see that despite the agent learning the value of ECS with time, this parameter still contributes the most to the variance of the SCC. To understand this result, we need to remember that we study the impact of the uncertainty around the climate feedback parameter at the initial state. This implies that even though the agent can learn about the value of the ECS over time, the current – year 2015 – range of possible prior means affects the variance of the SCC for the next hundred years.

Figures 8 and 9 present the univariate effects of the uncertain parameters on the value of the SCC for the models with and without learning, respectively. They depict the direction of the influence of the uncertain parameters on the absolute value of




Figure 6: Impact of the uncertain parameters on the SCC in the year 2020 (left), in 2050 (middle), and in 2100 (right) for the IAM with no learning, generated with 1,000 samples.




Figure 7: Impact of the uncertain parameters on the SCC in the year 2020 (left), in the year 2050 (middle), and in the 2100 (right) for the IAM with learning, generated with 1,000 samples.

the SCC. For all parameters, the direction of the impact corresponds to the economic intuition behind it: the SCC grows with an increasing prior mean of the feedback parameter, the damages, and IES. On the other side, the SCC decreases in the pure rate of time preferences. Finally, the SCC does not change in relative risk aversion and the prior variance of the climate feedback parameter.




Figure 8: Univariate effects on the SCC in the year 2020 for the model with no learning.




Figure 9: Univariate effects on the SCC in the year 2020 for the model with learning.

5.3 Tail learning and the social cost of carbon

In this section, we study how fat upper-tailed uncertainty over the temperature change from a doubling of GHG, that is, the ECS, affects economic policies. In addition, we examine whether and how fast uncertainties could be diminished through Bayesian learning (cf. Section 3.3).

Recall that climate change uncertainty, especially related to the ECS, poses significant challenges to policymaking. This uncertainty is characterized by a fat-tailed distribution, indicating a non-negligible probability of severe climate outcomes (see, e.g., Kelly and Tan (2015), Hwang et al. (2017), Fitzpatrick and Kelly (2017), and references therein). Consequently, an insurance motive, that is, protection against a possible adverse outcome, for reducing GHG emissions as a preventive measure against extreme temperature changes, emerges. However, climate policy is complicated by the fact that

learning about climate sensitivity progresses over time, influencing the urgency and extent of initial abatement actions.

Some of the previous literature (see, e.g., [Kelly and Kolstad \(1999\)](#)) suggests that learning about climate sensitivity, especially from a Bayesian perspective, is a slow process. This is due to factors such as stochastic weather events and uncertainties in climate feedback mechanisms, which obscure clear understanding ([Roe and Baker, 2007](#)). Despite this, rapid learning, or “tail learning,” might occur, leading to quicker rejection of severe climate scenarios ([Kelly and Tan, 2015](#)). This possibility challenges the common belief that resolving uncertainties in climate sensitivity is inherently slow.

From these stylized facts, the main question for climate policy revolves around the rate of resolving uncertainties and the optimal strategy in the presence of such uncertainties. If tail learning can be relatively fast, this has strong implications for near-term abatement policies. The presence of fat-tailed risk alters the optimal carbon tax and emissions levels initially, but learning can significantly reduce these adjustments over time.

To quantitatively address the implications of varying assumptions about the ECS, prior studies have typically resorted to executing a few computationally intensive simulations, each of which had a fixed parameterization, including the prior distribution, as they were suffering from the curse of dimensionality inflicted by the state variables as well as parameters. Our Deep UQ methodology offers a stark contrast to this approach. We leverage our surrogate model that incorporates also the prior mean and variance of the Bayesian learning process concerning the ECS. Recall that these parameters are directly added as pseudo-state variables in our model (cf. Table 1). This measure enables our IAM to cover a broad range of potential prior distributions, thereby facilitating a more efficient and comprehensive analysis of the impacts of varying ECS assumptions, as all required information can be obtained by interpolating the surrogate.²⁹

Following [Kelly and Tan \(2015\)](#), we define tail learning to be complete when the planner rejects ECS values that are located in the tail of the skewed distribution at a given confidence level. Formally, tail learning is complete at time period t if the planner rejects the hypothesis $\Delta T_{AT,\times 2,t} > \Delta T_{AT,\times 2,t}^L$, where $\Delta T_{AT,\times 2,t}^L$ is a lower bound of the tail distribution. There is no scientific consensus on the choice of $\Delta T_{AT,\times 2,t}^L$, therefore, following [Kelly and Tan \(2015\)](#), we set $\Delta T_{AT,\times 2,t}^L = \bar{\Delta T}_{AT,\times 2,t} + 1.5$, where $\bar{\Delta T}_{AT,\times 2,t}$ is estimated from the mean of the climate feedback parameter at time period t . Given the high damages caused by climate change, the planner requires a high level of confidence when he rejects the hypothesis at time period t . Thus, we consider the 99% and 99.9% confidence intervals on the ECS.

To examine how long the planner needs to complete the tail learning, we consider three particular values for the true ECS, namely $\Delta T_{AT,\times 2}^* = [2.0, 3.42, 4.5]$. For each of

²⁹ The role of fat-tailed uncertainty in near-term climate policy, particularly in the context of learning, is a subject of active debate in contemporary literature. [Weitzman \(2011\)](#), for instance, posits that due to the inherent inertia in climate systems and the challenge of rapidly reducing GHG emissions, learning about more severe climate change scenarios may have limited immediate policy relevance. In contrast, [Nordhaus \(2011\)](#) suggests that clear indicators of severe climate change will become apparent in the forthcoming 50 years, thereby providing a window for effective GHG reduction. [Pindyck \(2011\)](#), however, emphasizes that the answer to such questions is fundamental of a quantitative nature. Our framework, in turn, can support such discussions with quantitative answers.

those values, we record the first period t^* when the hypothesis $\Delta T_{AT,\times 2,t} > \Delta T_{AT,\times 2,t}^L$ is rejected and not subsequently not-rejected. We repeat this process $N^{\text{tail}} = 1,000$ times where the planner faces different realization of the climate feedback parameter \tilde{f} and the temperature shock ϵ_T , starting from the same prior of ECS. In Table 2, we report the expected tail learning time for each of the tree ECS under investigation. For the lowest of the three ECS values, learning is completed in less than a decade. For the higher values of the true ECS, this time increases as learning becomes progressively more difficult, since $\Delta T_{AT,\times 2}^*$ increases. However, even for $\Delta T_{AT,\times 2}^* = 4.5$, learning completes in only about 27 years at the 99% confidence level.

Table 2: Expected years until the tail learning is complete. The true ECS $\Delta T_{AT,\times 2}^* = 2.0$ is shown in the first row, $\Delta T_{AT,\times 2}^* = 3.42$ in the second row, and $\Delta T_{AT,\times 2}^* = 4.5$ is in the third row. The results are based on 1,000 simulated paths from 2015 to 2100.

	99%	99.9%
$\Delta T_{AT,\times 2}^* = 2.0$	5.46	9.06
$\Delta T_{AT,\times 2}^* = 3.42$	16.40	24.23
$\Delta T_{AT,\times 2}^* = 4.5$	27.26	38.57

Figure 10 depicts the evolution of the posterior probability density function of the ECS over time for an IAM with, that is, years 2015, 2050, 2100, and varying prior values. All other parameters were kept at their respective benchmark values (cf. Table 1, including $\psi = 1.5$).

Our results show that the social planner learns the upper tail of the prior distribution quickly. If the true climate sensitivity is moderate, that is, 3.42 or lower (cf. Figure 10a and Figure 10b), the tail learning is complete in about a few decades. These results confirm the findings of [Kelly and Tan \(2015\)](#), but contradict, for instance, [Weitzman \(2009\)](#) who claim that reducing uncertainty in the tail of the climate sensitivity prior to distribution must be a slow process since climate disasters are rare. In the case where the ECS is set to 4.5, that is, relatively high, learning slows (cf. Figure 10c). This is, as was also pointed out by [Kelly and Tan \(2015\)](#), due to the fact that the Bayes rule requires more observations to move the mean estimate from the prior ECS to the true high value.

Next, we quantify the effect of uncertainty on near-term optimal emissions, abatement policy, and the SCC. Table 3 presents the SCC for various true ECS values. Table 3 reveals that with an increasing value of the true ECS, both the mean as well as the SCC raises. This finding reflects the fact that more abatement is needed in case of more global warming, that is, higher temperatures. However, even in the case of the largest ECS value we consider Table 3, the SCC remains well below the levels of the SCC that are attainable in the model without learning (cf. Figure 4a).

To summarize, this section highlights the nuanced role of learning in climate policy, emphasizing the importance of understanding and adapting to evolving knowledge about the ECS. Furthermore, note that due to our Deeq UQ apparatus, the results presented in this section could be obtained in a computationally highly tractable fashion, that is, cheap interpolations on a surrogate rather than repeatedly computing and sim-

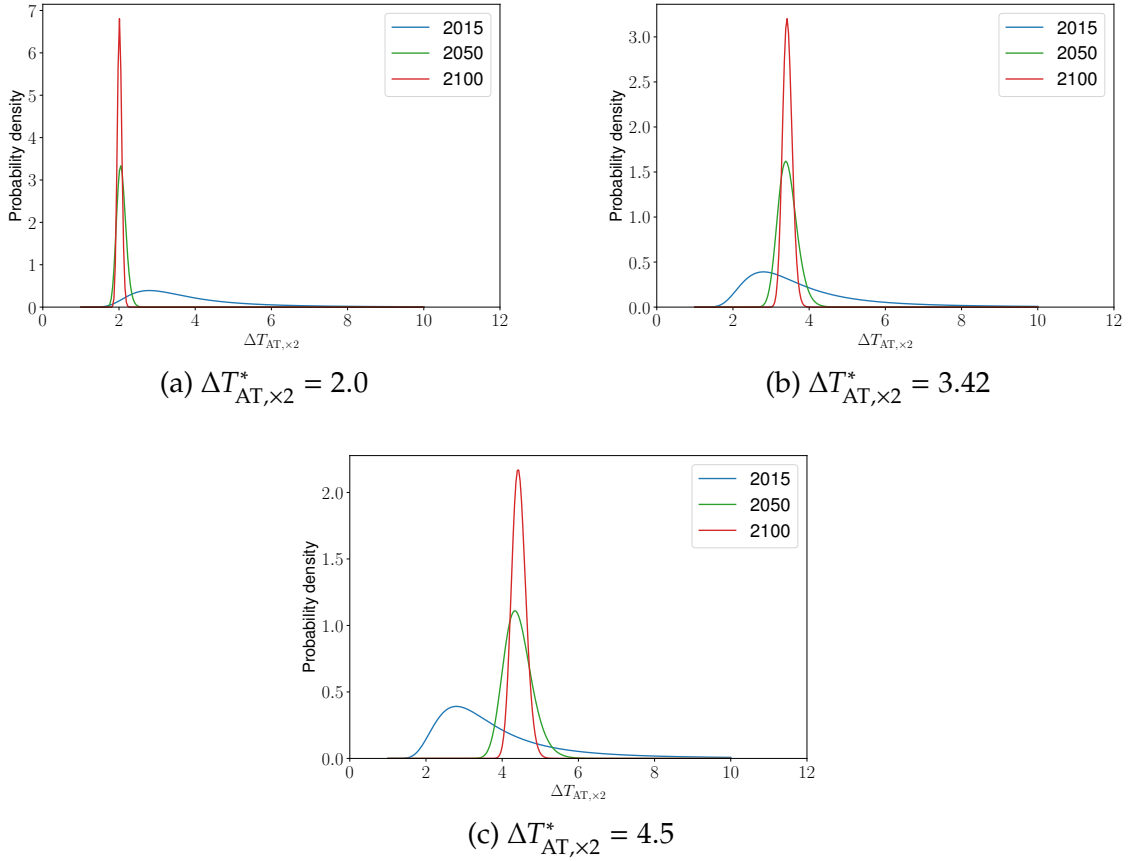


Figure 10: The posterior probability density function of $\Delta T_{AT,\times 2}$ is shown when the true ECS $\Delta T_{AT,\times 2}^*$ is set to 2.0, 3.42, and 4.5, respectively.

Table 3: SCC for the true ECS values $\Delta T_{AT,\times 2}^* = 2.0$ (left column), $\Delta T_{AT,\times 2}^* = 3.42$ (middle column), and $\Delta T_{AT,\times 2}^* = 4.5$ (right column), computed from 10,000 simulated paths for the years 2015 through 2100.

	$\Delta T_{AT,\times 2}^* = 2.0$	$\Delta T_{AT,\times 2}^* = 3.42$	$\Delta T_{AT,\times 2}^* = 4.5$
Mean in 2020	53.52	66.86	71.74
Mean in 2050	84.98	137.25	158.05
Mean in 2100	214.18	379.80	447.96
Standard deviation in 2020	5.28	5.62	5.73
Standard deviation in 2050	6.13	6.32	6.11
Standard deviation in 2100	10.59	10.15	11.15

ulating very expensive models for a vast set of parameters. We thereby show that we can solve very rich stochastic IAMs globally within minutes to hours and perform UQ on them on a laptop (cf. appendix B.2) rather than having to resort to supercomputers and use tens of thousands of nodes hours of compute time, as the previous literature had to do (see, for example, [Cai and Lontzek \(2019\)](#)).

6 Conclusion

This article posits a high-dimensional stochastic IAM that aligns with cutting-edge climate science. Our model incorporates a social planner with recursive preferences, iterative belief updates of equilibrium climate sensitivity using Bayes' rule, and stochastic climate tipping.

In response to the computational challenges present in this IAM, we introduce a generic methodological approach that consists of two main components: First, a deep learning-based algorithm designed to solve IAMs globally as a function of endogenous and exogenous state variables as well as uncertain parameters within a single model evaluation. Second, a Gaussian process-based surrogate model facilitates the efficient analysis of key UQ metrics, such as the SCC, with respect to uncertain model parameters. Our approach enables a rapid estimation of Sobol' indices, Shapley values, and univariate effects, which would otherwise be computationally very challenging.

Our numerical experiments indicate that the primary contributors to the variability in the SCC are parametric uncertainties in both the ECS and the damage function. Furthermore, our findings reveal that the uncertainty surrounding the ECS is largely resolved within approximately ten years. This resolution results in elevated optimal temperatures and a lower SCC, especially when compared to scenarios that do not incorporate Bayesian learning. Additionally, thanks to our Deeq UQ approach, we are able to study these effects through cost-effective interpolations using surrogate models, instead of the need for resource-intensive computations and simulations across an extensive parameter range. This advancement makes large-scale IAMs accessible to researchers who only have laptops at their disposal, eliminating the necessity for supercomputers.

In summary, our Deep UQ methodology enables a highly nuanced quantitative exploration of the intricate economic mechanisms at play in decision-making processes under uncertainty within IAMs. This approach opens up new pathways for further research in this field.

Appendix A Additional Model Details

This Appendix presents the complete parametrization of our annually calibrated IAM.

A.1 Exogenous variables

The law of motion for labor is given in Eq. (56) along with labor growth presented in Eq. (57), that is,

$$L_t = L_0 + (L_\infty - L_0) \left(1 - \exp \left(-\Delta_t \delta^L t \right) \right), \quad (56)$$

$$g_t^L = \frac{\frac{dL_t}{dt}}{L_t} = \frac{\Delta_t \delta^L}{\frac{L_\infty}{L_\infty - L_0} \exp \left(\Delta_t \delta^L t \right) - 1}. \quad (57)$$

The numerical values of the parameters for the world population and its growth rate are presented in Table 4:

Calibrated parameter	Symbol	Value
Annual rate of convergence	δ^L	0.0268
World population at starting year [millions]	L_0	7,403
Asymptotic world population [millions]	L_∞	11,500

Table 4: Parameterization for the evolution of labor.

The total factor productivity evolves according to Eq. (58), and the corresponding growth rate is presented in Eq. (59):

$$A_t = A_0 \exp \left(\frac{\Delta_t g_0^A (1 - \exp(-\Delta_t \delta^A t))}{\Delta_t \delta^A} \right). \quad (58)$$

$$g_t^A = \frac{\frac{dA_t}{dt}}{A_t} = \Delta_t g_0^A \exp \left(-\Delta_t \delta^A t \right). \quad (59)$$

The parameters of the total factor productivity evolution and TFP growth rate are provided in Table 5:

Calibrated parameter	Symbol	Value
The initial growth rate for TFP per year	g_0^A	0.0217
Decline rate of TFP growth per year	δ^A	0.005
Initial level of TFP	A_0	0.010295

Table 5: Parameterization for the evolution of TFP.

The carbon intensity is given by:

$$\sigma_t = \sigma_0 \exp\left(\frac{\Delta_t g_0^\sigma}{\log(1 + \Delta_t \delta^\sigma)} ((1 + \Delta_t \delta^\sigma)^t - 1)\right). \quad (60)$$

Furthermore, the parametrization for carbon intensity processes is given in Table 6:

Calibrated parameter	Symbol	Value
Initial growth of carbon intensity per year	g_0^σ	-0.0152
Decline rate of decarbonization per year	δ^σ	0.001
Initial carbon intensity (1000GtC)	σ_0	0.00009556

Table 6: Parameterization for the carbon intensity evolution.

Our IAM uses a backstop technology capable of mitigating the full amount of industrial emissions that enter the atmosphere. The cost of the backstop technology is assumed to be initially high but could be reduced over time, which is reflected in the definition of the coefficient of the abatement cost function $\theta_{1,t}$ as defined in Eq. (61).³⁰ The abatement cost is given by:

$$\theta_{1,t} = \frac{p_0^{\text{back}} \exp(-g^{\text{back}} t) 1000 \cdot \text{c2co2} \cdot \sigma_t}{\theta_2}. \quad (61)$$

The parameters for the abatement cost are given in Table 7:

Calibrated parameter	Symbol	Value
Cost of backstop 2010 thUSD per tCO ₂ 2015	p_0^{back}	0.55
Initial cost decline backstop cost per year	g^{back}	0.005
Exponent of control cost function	θ_2	2.6
Transformation coefficient from C to CO ₂	c2co2	3.666

Table 7: Parameterization for the abatement cost.

The non-industrial emissions from land use and deforestation decline over time according to Eq. (62), with a parameterization given in Table 8:

$$E_{\text{Land},t} = E_{\text{Land},0} \exp\left(-\Delta_t \delta^{\text{Land}} t\right). \quad (62)$$

³⁰ The scale parameter 1000 in Eq. (61) reflects the fact that we use a 1000 GtC unit of measurement; the parameter c2co2 transforms carbon intensity measured in GtC into GtCO₂, as the backstop price in DICE-2016 is given for GtCO₂ instead of GtC.

Calibrated parameter	Symbol	Value
Emissions from land 2015 (1000GtC per year)	$E_{\text{Land},0}$	0.000709
Decline rate of land emissions (per year)	δ^{Land}	0.023

Table 8: Parametrization for the emissions from land.

The exogenous radiative forcing resulting from non-CO₂ greenhouse gas emissions are described in Eq. (63):

$$F_t^{\text{EX}} = F_0^{\text{EX}} + \frac{1}{T/\Delta t} (F_1^{\text{EX}} - F_0^{\text{EX}}) \min(t, T/\Delta t). \quad (63)$$

The parameters of the exogenous radiative forcing are given in Table 9:

Calibrated parameter	Symbol	Value
2015 forcings of non-CO ₂ GHG (Wm-2)	F_0^{EX}	0.5
2100 forcings of non-CO ₂ GHG (Wm-2)	F_1^{EX}	1.0
Number of years before 2100	T	85

Table 9: Parametrization for the exogenous forcing.

A.2 The climate model

The parameters for the laws of motion for the masses of carbon, as well as the starting values and equilibrium values, are given in Table 10. We define: $b_{21} = b_{12} \frac{M_{\text{EQ}}^{\text{AT}}}{M_{\text{EQ}}^{\text{UO}}}$,

$$b_{32} = b_{23} \frac{M_{\text{EQ}}^{\text{UO}}}{M_{\text{EQ}}^{\text{LO}}}.$$

Calibrated parameter	Symbol	Value
Carbon cycle, annual value	b_{12}	0.054
Carbon cycle, annual value	b_{23}	0.0082
Equilibrium concentration in atmosphere (1000GtC)	$M_{\text{EQ}}^{\text{AT}}$	0.607
Equilibrium concentration in upper strata (1000GtC)	$M_{\text{EQ}}^{\text{UO}}$	0.489
Equilibrium concentration in lower strata (1000GtC)	$M_{\text{EQ}}^{\text{LO}}$	1.281
Concentration in atmosphere 2015 (1000GtC)	$M_{\text{INI}}^{\text{AT}}$	0.851
Concentration in upper strata 2015 (1000GtC)	$M_{\text{INI}}^{\text{UO}}$	0.628
Concentration in lower strata 2015 (1000GtC)	$M_{\text{INI}}^{\text{LO}}$	1.323

Table 10: Parametrization for the mass of carbon.

The parameters and starting values of temperature evolution and Bayesian learning are presented in Table 12:

Calibrated parameter	Symbol	Value
Temperature coefficient, annual value	c_1	0.137
Temperature coefficient, annual value	c_3	0.73
Temperature coefficient, annual value	c_4	0.00689
Forcings of equilibrium CO ₂ doubling (Wm ⁻²)	F_{2XCO_2}	3.45
Eq temperature impact (°C per doubling CO ₂)	T_{2XCO_2}	3.25
Eq concentration in atmosphere (1000GtC)	M_{base}^{AT}	0.607
Atmospheric temp change (°C) from 1850	T_0^{AT}	1.1
Lower stratum temp change (°C) from 1850	T_0^{OC}	0.27

Table 11: Parametrization for the temperature.

The parameters and starting values of Bayesian learning are given in Table 12:

Calibrated parameter	Symbol	Value
Initial prior mean	$\mu_{f,0}$	0.65
Initial prior variance	$S_{f,0}$	0.0169
Upper bound for a climate feedback parameters	\bar{f}	0.85
Lower bound for a climate feedback parameters	\underline{f}	0.45
Reference temperature impact (°C per doubling CO ₂)	$\bar{\Delta T}_{ATX2}^{\circ}$	1.2
Variance of temperature shock	S_{ϵ^T}	0.01

Table 12: Parametrization for the Bayesian learning process.

A.3 The economy

The baseline parameters for the economic side of the IAM can be found in Table 13.

Calibrated parameter	Symbol	Value
Capital annual depreciation rate	δ^K	0.1
Elasticity of capital	α	0.3
Damage parameter	ψ_1	20.46
Damage parameter	π_1	0.0
Damage parameter	π_2	0.00236
Tipping point variance	S_{TP}	0.01
Exponent of control cost function	θ_2	2.6
Intertemporal elasticity of substitution	ψ	1.5
Risk aversion	γ	10.
Time preferences	ρ	0.015

Table 13: Parameterization of the economic variables.

Appendix B Implementation Details

In this section, we discuss some practical points when an IAM is mapped to the DEQN solution method.

B.1 Deep equilibrium nets for stochastic IAMs

In this Appendix, we provide a detailed procedure for mapping a stochastic, nonlinear, and non-stationary IAM onto the DEQN solution framework. The original problem, as introduced in Section 3.4 (cf. Eqs. (21) to (32)), requires some modifications to leverage the capabilities of the DEQN.

Recall that the state of the economy at time t is given by

$$\mathbf{x}_t \in \mathbb{R}^{10+N} := \left(K_t, M_t^{\text{AT}}, M_t^{\text{UO}}, M_t^{\text{LO}}, T_t^{\text{AT}}, T_t^{\text{OC}}, \mu_{f,t}, S_{f,t}, TP_t, t, \boldsymbol{\vartheta} \right)^T, \quad (64)$$

where, aside from the endogenous and exogenous state variables, we consider N uncertain parameters as pseudo-state variables.

Following, for instance [Folini et al. \(2023\)](#), we consider time t as an exogenous state to account for the non-stationary nature of the IAM, whereas all the other states except the pseudo-states $\boldsymbol{\vartheta}$ are endogenously determined. Furthermore, to ensure computational tractability, we follow [Traeger \(2014\)](#) and map the unbounded physical time $t \in [0, \infty)$ via the strictly monotonic transformation,

$$\tau = 1 - \exp(-\zeta t), \quad (65)$$

into the unit interval $\tau \in (0, 1]$. To scale back from the artificial time τ to the physical time, the inverse transformation of Eq. (65) can be applied, that is,

$$t = -\frac{\ln(1 - \tau)}{\zeta}. \quad (66)$$

The dynamic programming problem presented in Section 3.4 could be computationally inefficient and unstable, mainly due to the capital stock that increases significantly over time. Thus, we follow [Traeger \(2014\)](#) and we normalize economic variables, that is, capital stock, consumption, and investment, as well as the value function in effective labor units, that is,

$$c_t := \frac{C_t}{A_t L_t}, k_t := \frac{K_t}{A_t L_t}, i_t := \frac{I_t}{A_t L_t}, v_t := \frac{V_t}{A_t L_t^{\frac{1}{1-\psi}}}, \quad (67)$$

where, as mentioned above, A_t represents a deterministic TFP growth trend and L_t is the global population (or labor). Furthermore, we introduce a quantity called the effective, or growth-adjusted, discount factor, that reads

$$\hat{\beta}_t := \exp \left(-\rho + \left(1 - \frac{1}{\psi} \right) g_t^A + g_t^L \right). \quad (68)$$

Using Eq. (67), we can transform the original dynamic programming problem with those quantities. Furthermore, we transform Eqs. (19) and (20) accordingly, and also leverage the fact that $\tilde{f}_{t+1} = \mu_{f,t} + \sqrt{S_{f,t}}\epsilon^f$, where $\epsilon^f \sim \mathcal{N}(0, 1, \underline{\epsilon}^f, \overline{\epsilon}^f)$, $\underline{\epsilon}^f = \frac{f - \mu_{f,t}}{\sqrt{S_{f,t}}}$, $\overline{\epsilon}^f = \frac{\bar{f} - \mu_{f,t}}{\sqrt{S_{f,t}}}$. For simplicity we replace $\epsilon_{T,t+1}$ from Eqs. (19) and (20) with by the shorthand notation ϵ^T . The laws of motion for posterior mean and posterior variance can be expressed as:

$$\mu_{f,t+1} = \mu_{f,t} + \frac{\sqrt{S_{f,t}}S_{f,t}(\varphi_{1C}T_{AT,t})^2}{S_{\epsilon^T} + (\varphi_{1C}T_{AT,t})^2 S_{f,t}}\epsilon^f + \frac{\varphi_{1C}T_{AT,t}S_{f,t}}{S_{\epsilon^T} + (\varphi_{1C}T_{AT,t})^2 S_{f,t}}\epsilon^T, \quad (69)$$

$$S_{f,t+1} = \frac{S_{\epsilon^T}S_{f,t}}{S_{\epsilon^T} + (\varphi_{1C}T_{AT,t})^2 S_{f,t}}, \quad (70)$$

where $\epsilon^f \sim \mathcal{N}(0, 1, \underline{\epsilon}^f, \overline{\epsilon}^f)$, $\underline{\epsilon}^f = \frac{f - \mu_{f,t}}{\sqrt{S_{f,t}}}$, $\overline{\epsilon}^f = \frac{\bar{f} - \mu_{f,t}}{\sqrt{S_{f,t}}}$ and $\epsilon^T \sim \mathcal{N}(0, S_{\epsilon^T})$.

The normalized recursive problem now reads as follows:

$$v_t(\mathbf{x}_t) = \max_{k_{t+1}, c_t, \mu_t} \left\{ \frac{c_t^{1-1/\psi}}{1-1/\psi} + \beta_t \mathbb{E}_t \left[v_{t+1}(\mathbf{x}_{t+1})^{\frac{1-\gamma}{1-\psi}} \right]^{\frac{1-1/\psi}{1-\gamma}} \right\} \quad (71)$$

$$\text{s.t.} \quad (1 - \Theta(\mu_t)) \Omega(T_{AT,t}) k_t^\alpha - c_t + (1 - \delta) k_t - \exp(g_t^A + g_t^L) k_{t+1} = 0 \quad (\lambda_t) \quad (72)$$

$$1 - \mu_t \geq 0 \quad \perp \quad \lambda_t^\mu \geq 0 \quad (73)$$

$$(1 - b_{12}) M_{AT,t} + b_{21} M_{UO,t} + (1 - \mu_t) \sigma_t A_t L_t k_t^\alpha + E_{Land,t} - M_{AT,t+1} = 0 \quad (v_t^{AT}) \quad (74)$$

$$b_{12} M_{AT,t} + (1 - b_{21} - b_{23}) M_{UO,t} + b_{32} M_{LO,t} - M_{UO,t+1} = 0 \quad (v_t^{UO}) \quad (75)$$

$$b_{23} M_{UO,t} + (1 - b_{32}) M_{LO,t} - M_{LO,t+1} = 0 \quad (v_t^{LO}) \quad (76)$$

$$(1 - c_1 c_3 - \varphi_{1C}) T_{AT,t} + c_1 c_3 T_{OC,t} + c_1 \left(F_{2\text{xc}02} \log_2 \left(\frac{M_{AT,t}}{M_{AT}^*} \right) + F_{EX,t} \right) \\ + \varphi_{1C} T_{AT,t} \mu_{f,t} + \varphi_{1C} T_{AT,t} \sqrt{S_{f,t}} \epsilon^f + \epsilon^T - T_{AT,t+1} = 0 \quad (\eta_{t+1}^{AT}), \quad (77)$$

$$\text{where } \epsilon^f \sim \mathcal{N}(0, 1, \underline{\epsilon^f}, \overline{\epsilon^f}), \epsilon^T \sim \mathcal{N}(0, S_{\epsilon^T})$$

$$c_4 T_{AT,t} + (1 - c_4) T_{OC,t} - T_{OC,t+1} = 0 \quad (\eta_t^{OC}) \quad (78)$$

$$\mu_{f,t} + \frac{\sqrt{S_{f,t}} S_{f,t} (\varphi_{1C} T_{AT,t})^2}{S_{\epsilon^T} + (\varphi_{1C} T_{AT,t})^2 S_{f,t}} \epsilon^f + \frac{\varphi_{1C} T_{AT,t} S_{f,t}}{S_{\epsilon^T} + (\varphi_{1C} T_{AT,t})^2 S_{f,t}} \epsilon^T - \mu_{f,t+1} = 0 \quad (\lambda_{t+1}^{\mu_f})$$

$$\text{where } \epsilon^f \sim \mathcal{N}(0, 1, \underline{\epsilon^f}, \overline{\epsilon^f}), \epsilon^T \sim \mathcal{N}(0, S_{\epsilon^T}) \quad (79)$$

$$\frac{S_{\epsilon^T} S_{f,t}}{S_{\epsilon^T} + (\varphi_{1C} T_{AT,t})^2 S_{f,t}} - S_{f,t+1} = 0 \quad (\lambda_t^{S_f}) \quad (80)$$

$$\Omega(T_{AT,t}) = 1 - \frac{1}{1 + \left(\frac{1}{20.46} T_{AT,t}\right)^2 + \left(\frac{1}{2 \cdot TP_t} T_{AT,t}\right)^{6.754}} \quad (81)$$

$$TP_{t+1} = TP_t + \epsilon_{TP,t+1}, \quad \epsilon_{TP,t} \sim \mathcal{N}(0, S_{TP}) \quad (82)$$

where the Lagrange multipliers we will employ below have been added in parentheses for completeness. In Eq. (73), the symbol \perp indicates complementary slackness.

The policy function \mathbf{p} we intend to approximate with the aid of deep neural networks is given by

$$\mathcal{N}(\mathbf{x}_t) \in \mathbb{R}^{11} := \left(k_{t+1}, c_t, \mu_t, v_t^{AT}, v_t^{UO}, v_t^{LO}, \eta_{t+1}^{AT}, \eta_t^{OC}, \lambda_{t+1}^{\mu_f}, \lambda_t^{S_f}, v_t \right), \quad (83)$$

and consists of the choice variables (k_{t+1}, c_t, μ_t) ³¹ as well as the Lagrange multipliers

³¹ We keep the first-order conditions with respect to k_{t+1} and c_t . In practical applications, the performance of neural networks can, as in our case, often benefit from redundant information (see, e.g.,

and the value function.

Next, we derive the first-order conditions to form a loss function for the IAM model that is suitable for DEQNs (cf. Eq. (37)) in effective labor units.

$$\frac{\partial v_t}{\partial k_{t+1}} = \beta_t \mathbb{E}_t \left[v_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right]^{\frac{\gamma-1/\psi}{1-\gamma}} \mathbb{E}_t \left[v_{t+1}^{\frac{1/\psi-\gamma}{1-1/\psi}} v_{k,t+1} \right] - \lambda_t \exp(g_t^A + g_t^L) = 0 \quad (84)$$

$$\frac{\partial v_t}{\partial c_t} = c_t^{-\frac{1}{\psi}} - \lambda_t = 0 \quad (85)$$

$$\frac{\partial v_t}{\partial \mu_t} = \lambda_t \Theta'(\mu_t) \Omega_t(T_{AT,t}) k_t^\alpha + \lambda_t^\mu + v_t^{AT} \sigma_t A_t L_t k_t^\alpha = 0 \quad (86)$$

$$\frac{\partial v_t}{\partial M_{AT,t+1}} = \beta_t \mathbb{E}_t \left[v_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right]^{\frac{\gamma-1/\psi}{1-\gamma}} \mathbb{E}_t \left[v_{t+1}^{\frac{1/\psi-\gamma}{1-1/\psi}} v_{M_{AT},t+1} \right] - v_t^{AT} = 0 \quad (87)$$

$$\frac{\partial v_t}{\partial M_{UO,t+1}} = \beta_t \mathbb{E}_t \left[v_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right]^{\frac{\gamma-1/\psi}{1-\gamma}} \mathbb{E}_t \left[v_{t+1}^{\frac{1/\psi-\gamma}{1-1/\psi}} v_{M_{UO},t+1} \right] - v_t^{UO} = 0 \quad (88)$$

$$\frac{\partial v_t}{\partial M_{LO,t+1}} = \beta_t \mathbb{E}_t \left[v_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right]^{\frac{\gamma-1/\psi}{1-\gamma}} \mathbb{E}_t \left[v_{t+1}^{\frac{1/\psi-\gamma}{1-1/\psi}} v_{M_{LO},t+1} \right] - v_t^{LO} = 0 \quad (89)$$

$$\begin{aligned} \frac{\partial v_t}{\partial T_{AT,t+1}} &= \beta_t \mathbb{E}_t \left[v_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right]^{\frac{\gamma-1/\psi}{1-\gamma}} \mathbb{E}_t \left[v_{t+1}^{\frac{1/\psi-\gamma}{1-1/\psi}} v_{T_{AT},t+1} \right] \\ &- \int_{-\infty}^{+\infty} \int_{\underline{\epsilon^f}}^{\overline{\epsilon^f}} \eta_{t+1}^{AT}(\epsilon^f, \epsilon^T) \text{pdf}(\epsilon^f) d\epsilon^f \text{pdf}(\epsilon^T) d\epsilon^T = 0 \end{aligned} \quad (90)$$

$$\frac{\partial v_t}{\partial T_{OC,t+1}} = \beta_t \mathbb{E}_t \left[v_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right]^{\frac{\gamma-1/\psi}{1-\gamma}} \mathbb{E}_t \left[v_{t+1}^{1/\psi-\gamma} 1 - 1/\psi v_{T_{OC},t+1} \right] - \eta_t^{OC} = 0 \quad (91)$$

$$\begin{aligned} \frac{\partial v_t}{\partial \mu_{f,t+1}} &= \beta_t \mathbb{E}_t \left[v_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right]^{\frac{\gamma-1/\psi}{1-\gamma}} \mathbb{E}_t \left[v_{t+1}^{\frac{1/\psi-\gamma}{1-1/\psi}} v_{\mu_f,t+1} \right] \\ &- \int_{-\infty}^{+\infty} \int_{\underline{\epsilon^f}}^{\overline{\epsilon^f}} \lambda_{t+1}^{\mu_f}(\epsilon^f, \epsilon^T) \text{pdf}(\epsilon^f) d\epsilon^f \text{pdf}(\epsilon^T) d\epsilon^T = 0 \end{aligned} \quad (92)$$

$$\frac{\partial v_t}{\partial S_{f,t+1}} = \beta_t \mathbb{E}_t \left[v_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right]^{\frac{\gamma-1/\psi}{1-\gamma}} \mathbb{E}_t \left[v_{t+1}^{\frac{1/\psi-\gamma}{1-1/\psi}} v_{S_f,t+1} \right] - \lambda_t^{S_f} = 0 \quad (93)$$

In equilibrium, the conditions following from the envelop theorem hold:

$$\begin{aligned} v_{k,t} &= \frac{\partial v_t}{\partial k_t} = \lambda_t \frac{\partial k_{t+1}}{\partial k_t} + v_t^{AT} \frac{\partial M_{AT,t+1}}{\partial k_t} \\ &\Leftrightarrow v_{k,t} = \lambda_t \left\{ (1 - \Theta(\mu_t)) \Omega(T_{AT,t}) \alpha k_t^{\alpha-1} + (1 - \delta) \right\} \\ &\quad + v_t^{AT} (1 - \mu_t) \sigma_t A_t L_t \alpha k_t^{\alpha-1} \end{aligned} \quad (94)$$

[Azinovic et al. \(2022\)](#) for more details).

$$\begin{aligned}
v_{M_{AT,t}} &= \frac{\partial v_t}{\partial M_{AT,t}} = v_t^{AT} \frac{\partial M_{AT,t+1}}{\partial M_{AT,t}} + v_t^{UO} \frac{\partial M_{UO,t+1}}{\partial M_{AT,t}} + \eta_{t+1}^{AT} \frac{\partial T_{AT,t+1}}{\partial M_{AT,t}} \\
&\Leftrightarrow v_{M_{AT,t}} = v_t^{AT} (1 - b_{12}) + v_t^{UO} b_{12} + \\
&c_1 F_{2 \times c_0} \frac{1}{\log 2 M_{AT,t}} \int_{-\infty}^{+\infty} \int_{\underline{\epsilon^f}}^{\overline{\epsilon^f}} \eta_{t+1}^{AT}(\epsilon^f, \epsilon^T) \text{pdf}(\epsilon^f) d\epsilon^f \text{pdf}(\epsilon^T) d\epsilon^T \quad (95)
\end{aligned}$$

$$\begin{aligned}
v_{M_{UO,t}} &= \frac{\partial v_t}{\partial M_{UO,t}} = v_t^{AT} \frac{\partial M_{AT,t+1}}{\partial M_{UO,t}} + v_t^{UO} \frac{\partial M_{UO,t+1}}{\partial M_{UO,t}} + v_t^{LO} \frac{\partial M_{LO,t+1}}{\partial M_{UO,t}} \\
&\Leftrightarrow v_{M_{UO,t}} = v_t^{AT} b_{21} + v_t^{UO} (1 - b_{21} - b_{23}) + v_t^{LO} b_{23} \quad (96)
\end{aligned}$$

$$\begin{aligned}
v_{M_{LO,t}} &= \frac{\partial v_t}{\partial M_{LO,t}} = v_t^{UO} \frac{\partial M_{UO,t+1}}{\partial M_{LO,t}} + v_t^{LO} \frac{\partial M_{LO,t+1}}{\partial M_{LO,t}} \\
&\Leftrightarrow v_{M_{LO,t}} = v_t^{UO} b_{32} + v_t^{LO} (1 - b_{32}) \quad (97)
\end{aligned}$$

$$\begin{aligned}
v_{T_{AT,t}} &= \frac{\partial v_t}{\partial T_{AT,t}} = \lambda_t \frac{\partial k_{t+1}}{\partial T_{AT,t}} + \eta_{t+1}^{AT} \frac{\partial T_{AT,t+1}}{\partial T_{AT,t}} + \eta_t^{OC} \frac{\partial T_{OC,t+1}}{\partial T_{AT,t}} + \lambda_{t+1}^{\mu_f} \frac{\partial \mu_{f,t+1}}{\partial T_{AT,t}} + \lambda_t^{S_f} \frac{\partial S_{f,t+1}}{\partial T_{AT,t}} \\
&\Leftrightarrow v_{T_{AT,t}} = \lambda_t (1 - \Theta(\mu_t)) \Omega'(T_{AT,t}) k_t^\alpha \\
&+ \int_{-\infty}^{+\infty} \int_{\underline{\epsilon^f}}^{\overline{\epsilon^f}} \eta_{t+1}^{AT}(\epsilon^f, \epsilon^T) \left((1 - c_1 c_3 - \varphi_{1C}) + \varphi_{1C} \mu_{f,t} + \varphi_{1C} \sqrt{S_{f,t}} \epsilon^f \right) \text{pdf}(\epsilon^f) d\epsilon^f \text{pdf}(\epsilon^T) d\epsilon^T \\
&+ \eta_t^{OC} c_4 + \\
&\int_{-\infty}^{+\infty} \int_{\underline{\epsilon^f}}^{\overline{\epsilon^f}} \lambda_{t+1}^{\mu_f}(\epsilon^f, \epsilon^T) \frac{2S_{f,t}^{3/2} \varphi_{1C}^2 T_{AT,t} S_{\epsilon^T} \epsilon^f + \varphi_{1C} S_{f,t} S_{\epsilon^T} \epsilon^T - \varphi_{1C}^3 T_{AT,t}^2 S_{f,t}^2 \epsilon^T}{(S_{\epsilon^T} + (\varphi_{1C} T_{AT,t})^2 S_{f,t})^2} \text{pdf}(\epsilon^f) d\epsilon^f \text{pdf}(\epsilon^T) d\epsilon^T \\
&- \lambda_t^{S_f} \frac{2S_{\epsilon^T} \varphi_{1C}^2 T_{AT,t} S_{f,t}^2}{(S_{\epsilon^T} + (\varphi_{1C} T_{AT,t})^2 S_{f,t})^2} \quad (98)
\end{aligned}$$

$$\begin{aligned}
v_{T_{OC,t}} &= \frac{\partial v_t}{\partial T_{OC,t}} = \eta_{t+1}^{AT} \frac{\partial T_{AT,t+1}}{\partial T_{OC,t}} + \eta_t^{OC} \frac{\partial T_{OC,t+1}}{\partial T_{OC,t}} \\
&\Leftrightarrow v_{T_{OC,t}} = c_1 c_3 \int_{-\infty}^{+\infty} \int_{\underline{\epsilon^f}}^{\overline{\epsilon^f}} \eta_{t+1}^{AT}(\epsilon^f, \epsilon^T) \text{pdf}(\epsilon^f) d\epsilon^f \text{pdf}(\epsilon^T) d\epsilon^T + \eta_t^{OC} (1 - c_4) \quad (99)
\end{aligned}$$

$$\begin{aligned}
v_{\mu_{f,t}} &= \frac{\partial v_t}{\partial \mu_{f,t}} = \eta_{t+1}^{AT} \frac{\partial T_{AT,t+1}}{\partial \mu_{f,t}} + \lambda_{t+1}^{\mu_f} \frac{\partial \mu_{f,t+1}}{\partial \mu_{f,t}} \\
&\Leftrightarrow v_{\mu_{f,t}} = \varphi_{1C} T_{AT,t} \int_{-\infty}^{+\infty} \int_{\underline{\epsilon^f}}^{\overline{\epsilon^f}} \eta_{t+1}^{AT}(\epsilon^f, \epsilon^T) \text{pdf}(\epsilon^f) d\epsilon^f \text{pdf}(\epsilon^T) d\epsilon^T + \\
&\int_{-\infty}^{+\infty} \int_{\underline{\epsilon^f}}^{\overline{\epsilon^f}} \lambda_{t+1}^{\mu_f}(\epsilon^f, \epsilon^T) \text{pdf}(\epsilon^f) d\epsilon^f \text{pdf}(\epsilon^T) d\epsilon^T \quad (100)
\end{aligned}$$

$$v_{S_{f,t}} = \frac{\partial v_t}{\partial S_{f,t}} = \eta_{t+1}^{AT} \frac{\partial T_{AT,t+1}}{\partial S_{f,t}} + \lambda_{t+1}^{\mu_f} \frac{\partial \mu_{f,t+1}}{\partial S_{f,t}} + \lambda_t^{S_f} \frac{\partial S_{f,t+1}}{\partial S_{f,t}}$$

$$\begin{aligned}
\Leftrightarrow v_{S_{f,t}} &= \int_{-\infty}^{+\infty} \int_{\underline{\epsilon^f}}^{\overline{\epsilon^f}} \eta_{t+1}^{\text{AT}}(\epsilon^f, \epsilon^T) \frac{1}{2} \varphi_{1C} T_{\text{AT},t} S_{f,t}^{-1/2} \epsilon^f \text{pdf}(\epsilon^f) d\epsilon^f \text{pdf}(\epsilon^T) d\epsilon^T \\
&+ \int_{-\infty}^{+\infty} \int_{\underline{\epsilon^f}}^{\overline{\epsilon^f}} \lambda_{t+1}^{\mu_f}(\epsilon^f, \epsilon^T) \\
&\frac{\frac{3}{2} (\varphi_{1C} T_{\text{AT},t})^2 \epsilon^f S_{\epsilon^T} + \frac{1}{2} S_{f,t} (\varphi_{1C} T_{\text{AT},t})^4 \epsilon^f + \varphi_{1C} T_{\text{AT},t} S_{\epsilon^T} \epsilon^T}{\left(S_{\epsilon^T} + (\varphi_{1C} T_{\text{AT},t})^2 S_{f,t} \right)^2} \text{pdf}(\epsilon^f) d\epsilon^f \text{pdf}(\epsilon^T) d\epsilon^T \\
&+ \lambda_t^{S_f} \frac{S_{\epsilon^T}^2}{\left(S_{\epsilon^T} + (\varphi_{1C} T_{\text{AT},t})^2 S_{f,t} \right)^2}
\end{aligned} \tag{101}$$

Finally, the following optimality condition holds:

$$v^* = \frac{c_t^{1-1/\psi}}{1-1/\psi} + \beta_t \mathbb{E}_t \left[v^* \frac{1-\gamma}{1-1/\psi} \right]. \tag{102}$$

and budget constraint reads as

$$(1 - \Theta(\mu_t)) \Omega(T_{\text{AT},t}) k_t^\alpha - c_t + (1 - \delta) k_t - \exp\left(g_t^A + g_t^L\right) k_{t+1} = 0. \tag{103}$$

We replace the KKT condition in Eq. (73) with the Fischer-Burmeister function (see, e.g., [Maliar et al. \(2021\)](#), and references therein) and directly embed it in the system of non-linear equilibrium conditions, that is,

$$\Psi^{\text{FB}}\left(\lambda_t^\mu, 1 - \mu_t\right) = \lambda_t^\mu + (1 - \mu_t) - \sqrt{\lambda_t^{\mu^2} + (1 - \mu_t)^2}, \tag{104}$$

where from Eq. (86), we define λ_t^μ such that

$$\lambda_t^\mu \equiv -\lambda_t \Theta'(\mu_t) \Omega(T_{\text{AT},t}) k_t^\alpha - v_t^{\text{AT}} \sigma_t A_t L_t k_t^\alpha. \tag{105}$$

One issue that arises when working with the first-order conditions of an IAM is that they need to be computed not only with respect to the economic choice variables, such as μ_t and c_t , but also with respect to the climate variables, even though they are not choice variables. The reason for this is that the marginal effects of the change in choice variables that propagate through the climate system need to be assessed. These effects cannot be computed analytically here, which is why we need Lagrange multipliers associated with every single climate equation (cf. Eqs. (87) to (91)) to estimate the shadow price of a marginal change in a respective constraint.

Using all the above definitions, the eleven individual components that enter the loss function amendable for the DEQN algorithm read as follows:

$$l_1 := \frac{c_t^{1-1/\psi}}{1-1/\psi} + \beta_t \mathbb{E}_t \left[v_{t+1} \frac{1-\gamma}{1-1/\psi} \right] - v_t \tag{106}$$

$$l_2 := (1 - \Theta(\mu_t)) \Omega(T_{\text{AT},t}) k_t^\alpha - c_t + (1 - \delta) k_t - \exp(g_t^A + g_t^L) k_{t+1} \quad (107)$$

$$l_3 := \lambda_t^\mu + (1 - \mu_t) - \sqrt{\lambda_t^{\mu^2} + (1 - \mu_t)^2} \quad (108)$$

$$l_4 := \lambda_t \exp(g_t^A + g_t^L) - \beta_t \mathbb{E}_t \left[v_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right]^{\frac{\gamma-1/\psi}{1-\gamma}} \mathbb{E}_t \left[v_{t+1}^{\frac{1/\psi-\gamma}{1-1/\psi}} v_{k,t+1} \right] \quad (109)$$

$$l_5 := v_t^{\text{AT}} - \beta_t \mathbb{E}_t \left[v_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right]^{\frac{\gamma-1/\psi}{1-\gamma}} \mathbb{E}_t \left[v_{t+1}^{\frac{1/\psi-\gamma}{1-1/\psi}} v_{M_{\text{AT}},t+1} \right] \quad (110)$$

$$l_6 := v_t^{\text{UO}} - \beta_t \mathbb{E}_t \left[v_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right]^{\frac{\gamma-1/\psi}{1-\gamma}} \mathbb{E}_t \left[v_{t+1}^{\frac{1/\psi-\gamma}{1-1/\psi}} v_{M_{\text{UO}},t+1} \right] \quad (111)$$

$$l_7 := v_t^{\text{LO}} - \beta_t \mathbb{E}_t \left[v_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right]^{\frac{\gamma-1/\psi}{1-\gamma}} \mathbb{E}_t \left[v_{t+1}^{\frac{1/\psi-\gamma}{1-1/\psi}} v_{M_{\text{LO}},t+1} \right] \quad (112)$$

$$l_8 := \int_{-\infty}^{+\infty} \int_{\underline{\epsilon^f}}^{\bar{\epsilon^f}} \eta_{t+1}^{\text{AT}}(\epsilon^f, \epsilon^T) \text{pdf}(\epsilon^f) d\epsilon^f \text{pdf}(\epsilon^T) d\epsilon^T \\ - \beta_t \mathbb{E}_t \left[v_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right]^{\frac{\gamma-1/\psi}{1-\gamma}} \mathbb{E}_t \left[v_{t+1}^{\frac{1/\psi-\gamma}{1-1/\psi}} v_{T_{\text{AT}},t+1} \right] \quad (113)$$

$$l_9 := \eta_t^{\text{OC}} - \beta_t \mathbb{E}_t \left[v_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right]^{\frac{\gamma-1/\psi}{1-\gamma}} \mathbb{E}_t \left[v_{t+1}^{\frac{1/\psi-\gamma}{1-1/\psi}} v_{T_{\text{OC}},t+1} \right] \quad (114)$$

$$l_{10} := \int_{-\infty}^{+\infty} \int_{\underline{\epsilon^f}}^{\bar{\epsilon^f}} \lambda_{t+1}^{\mu_f}(\epsilon^f, \epsilon^T) \text{pdf}(\epsilon^f) d\epsilon^f \text{pdf}(\epsilon^T) d\epsilon^T \\ - \beta_t \mathbb{E}_t \left[v_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right]^{\frac{\gamma-1/\psi}{1-\gamma}} \mathbb{E}_t \left[v_{t+1}^{\frac{1/\psi-\gamma}{1-1/\psi}} v_{\mu_f,t+1} \right] \quad (115)$$

$$l_{11} := \lambda_t^{S_f} - \beta_t \mathbb{E}_t \left[v_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right]^{\frac{\gamma-1/\psi}{1-\gamma}} \mathbb{E}_t \left[v_{t+1}^{\frac{1/\psi-\gamma}{1-1/\psi}} v_{S_f,t+1} \right] \quad (116)$$

and result in the total loss function given by

$$\ell_\nu := \frac{1}{N_{\text{path length}}} \sum_{\mathbf{x}_t \text{ on sim. path}} \sum_{m=1}^{N_{\text{eq}}=11} (l_m(\mathbf{x}_t, \mathcal{N}(\mathbf{x}_t)))^2. \quad (117)$$

The final ingredient we need for the DEQN algorithm is the evolution of the state \mathbf{x}_t one period forward such that the loss function (117) can be evaluated along a simulated path. In our application, \mathbf{x}_{t+1} is given by

$$\mathbf{x}_{t+1} = \left(k_{t+1}, M_{t+1}^{\text{AT}}, M_{t+1}^{\text{UO}}, M_{t+1}^{\text{LO}}, T_{t+1}^{\text{AT}}, T_{t+1}^{\text{OC}}, \mu_{f,t+1}, S_{f,t+1}, TP_{t+1}, t+1, \vartheta \right)^T, \quad (118)$$

where k_{t+1} is updated through the law of motion using a choice variable from the

policy function Eq. (83) and the climate variables M_{t+1}^{AT} , M_{t+1}^{UO} , M_{t+1}^{LO} , T_{t+1}^{AT} , and T_{t+1}^{OC} can be updated via Eqs. (74) to (80), tipping points follow an exogenous process Eq. (33) whereas time t is simply incremented by one unit, and the pseudo-states ϑ are re-sampled from their distribution at every iteration step.

B.2 Hyperparameter selection and neural network training

Solving the stochastic IAM under consideration with DEQN can be a challenging procedure due to the extensive hyperparameter space available for modelers. This process, therefore, requires some degree of experimentation. In our numerical experiments, we employed a neural network architecture with two hidden layers, each consisting of 1,000 neurons activated by the scaled exponential linear unit (SELU) function. We utilized the Adam optimizer (Kingma and Ba, 2014) with a learning rate of $\alpha^{\text{learn}} = 1 \cdot 10^{-5}$. Our approach involved choosing 512 points and simulating them forward. To stabilize training, we initially limited the simulation path length to $N_{\text{path}} = 4$, subsequently increasing it to 8, 16, 32, 64, 128, 256, and finally 300 steps, contingent upon the stabilization of the policies. Additionally, we began with a mini-batch size of 64, adjusting it to 128 towards the end of training when the training error was already low and exhibited negligible improvement.

Solving one IAM outlined in full detail in appendix B.1 from scratch with the settings just discussed until full convergence requires about 4 hours on an ordinary laptop. This time can be reduced to minutes if a pre-trained model solution exists. All the models presented in the paper were solved on an 8-core Intel compute node on <https://nuvolos.cloud> with 64GB of RAM, and 100GB of fast local storage (SSD).

Appendix C More details on Gaussian processes

C.1 Leave-one-out error with Gaussian processes

The selection of the number of sample points n contained in the training set \mathcal{D} is a trade-off between accuracy and efficiency. Our design philosophy aims to minimize n but still ensure the accuracy of the surrogate model predictions based on GPs in a computationally affordable way.

One common choice in the UQ literature to strike a balance between these two opposing factors, although primarily discussed with polynomial chaos expansions (see, e.g., Blatman and Sudret (2010), Le Gratiet et al. (2017), Harenberg et al. (2019)), is to use the *leave-one-out* (LOO) error estimator. Recall that, as in Eq. (45), we have a computational model $\mathcal{M}(\cdot)$ that maps an input vector x_i to a scalar output y_i . We repeatedly evaluate Eq. (45) n times to obtain a training dataset $\mathcal{D} = \{x_i, y_i\}_{i=1}^n = [X, y]$. Then, we fit the GP to the original model $\mathcal{M}(\cdot)$, denoted as $\mathcal{M}_{\text{GP}|X,y}(\cdot)$, as discussed in Section 4.1.2. To measure the LOO error, we first construct a GP surrogate model $\mathcal{M}_{\text{GP}|X_{-i},y_{-i}}$ on experimental design points $X_{-i} \equiv X \setminus x_i = \{x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n\}$, and estimate the error Δ_i on the excluded point x_i between the outcome from the true model and the prediction from the GP model, that is,

$$\Delta_i \equiv \mathcal{M}(x_i) - \mathcal{M}_{\text{GP}|X_{-i},y_{-i}}(x_i) \quad (119)$$

Next, we compute the sum of Δ_i over i and define the LOO error, that is,

$$\epsilon_{\text{LOO}}^{\text{GP}} \equiv \frac{1}{n} \sum_{i=1}^n \Delta_i^2 = \frac{1}{N} \sum_{i=1}^n (\mathcal{M}(x_i) - \mathcal{M}_{\text{GP}|X_{-i}, Y_{-i}}(x_i))^2. \quad (120)$$

In our numerical applications below, we choose the size of the training set n for the GP surrogate such that we achieve $\epsilon_{\text{LOO}}^{\text{GP}} \leq 10^{-2}$, which guarantees the required high accuracy of the GP surrogate for our QoIs, obtained at moderate computational costs. If a given initial n is insufficient, we systematically increase the size of the training set, for example, by applying Bayesian active learning, as will be discussed in appendix C.2.

C.2 Experimental design

To enhance the precision of our GP-based surrogate models in UQ experiments, the efficient selection of training observations for the dataset \mathcal{D} , used to fit the GPs, is crucial. As discussed in Section 4.1.2, this selection process, known as “experimental design,” systematically identifies a subset of input data points from the available space $\mathcal{B} \subset \mathbf{R}^N$. A comprehensive discussion on experimental design is available in [Santner et al. \(2018\)](#). In our target application, as is common in science, where stochastic simulations are employed, the generation of training data (e.g., for the SCC in 2100) relies on lengthy simulations, rendering it computationally expensive. Therefore, an accurate estimation of the relationship between model inputs and outputs with a limited number of simulations becomes imperative.

In the realm of GPR, the primary objective is to judiciously select the experimental design. This selection aims to minimize computational costs and uncertainty while ensuring accurate modeling of the target function or process. In essence, the challenge is to develop a training set that enables the creation of high-quality surrogates via GPs, using the least amount of data points feasible. Techniques like Bayesian optimization and Bayesian active learning (see, e.g., [Renner and Scheidegger \(2018\)](#), and references therein) are pivotal in this context. They facilitate the selection of data points that either maximize predictive performance or minimize prediction uncertainty. This approach addresses the “exploration-exploitation dilemma”, which involves balancing the acquisition of new statistical knowledge (exploration) against utilizing existing insights about an unknown function we seek to approximate with minimal observations (exploitation). Our goal is to improve the global accuracy of model predictions, focusing primarily on exploration. Subsequently, we will elaborate on the implementation of an efficient, model-dependent experimental design for our applications and compare the efficacy of three different sampling methods.

Given a set of parameter vectors ϑ_i within a compact set $[\underline{\vartheta}, \bar{\vartheta}] \subset \mathcal{B}$, $\underline{\vartheta}, \bar{\vartheta} \in \mathbf{R}^N$, as detailed in Table 1, the initial approach for generating training data involves adopting space-filling designs such as the Sobol sequence ([Sobol, 1967](#)), the Halton sequence ([Halton, 1960](#)), and Latin hypercube sampling (LHS; [McKay et al. \(1979\)](#)). LHS is particularly prominent in GSA ([Butler et al., 2014](#), [Harenberg et al., 2019](#)). However, [Chen and Zhou \(2014\)](#) noted that while space-filling designs are adequate for deterministic simulations, they may fall short for stochastic models. This shortfall arises because

these designs not only require that the locations of design points be pre-determined, but these designs also require that substantial computational resources be allocated to each point. Furthermore, non-sequential designs may be inefficient, especially when constructing surrogate models for high-dimensional input spaces.

An alternative approach involves Bayesian active learning with various types of “acquisition functions”. Recall that GPs, at each design point, provide both the posterior mean and variance with minimal computational effort (cf. Eqs. (41) and (42)). This feature is leveraged to sequentially query new candidate points, not initially included in the design, which most significantly enhance the model’s approximation quality. Such a strategy effectively minimizes the computational budget.³²

In our study, we adopt two design criteria for GP surrogate models as outlined by [Sacks et al. \(1989\)](#). The first criterion utilizes an acquisition function, denoted as α_{MSE} , which aims to minimize the Mean Squared Error (MSE). Following [Binois et al. \(2019\)](#), we use the noise-free posterior variance for computational simplicity. This MSE corresponds to the de-noised posterior variance:

$$\alpha_{\text{MSE}}(x) = k(x, x) - k(x, X)K(X, X)^{-1}k(X, x). \quad (121)$$

To enhance the basic training set, we select a point x_{n+1} from the domain of interest \mathcal{B} by solving the following optimization problem:

$$x_{n+1} = \arg \max_{x \in \mathcal{B}} \alpha_{\text{MSE}}(x), \quad (122)$$

and then add this point to the current set \mathcal{D} .

The second criterion we consider in our work is the Integrated Mean Square Error (IMSE). For this, we define α_{IMSE} according to the existing literature:

$$\alpha_{\text{IMSE}} = \int_{x \in \mathcal{B}} \left(k(x, x) - k(x, X)K(X, X)^{-1}k(X, x) \right) dx, \quad (123)$$

where the integration is performed over the domain of interest \mathcal{B} . Again we assume the de-noised posterior variance. The next point x_{n+1} is selected to minimize the IMSE:

$$x_{n+1} = \arg \min_{x \in \mathcal{B}} \alpha_{\text{IMSE}}(x). \quad (124)$$

These acquisition functions, defined in Eq. (121) and Eq. (123), are designed to improve the global predictive accuracy of a GP surrogate model by leveraging posterior variance information. Choosing between these criteria depends on the specific models under consideration. We further investigate the performance of the MSE and IMSE acquisition functions compared to a space-filling LHS design, examining their impact on the convergence of the LOO error.

We investigate the convergence of the LOO error for the SCC in the year 2100 employing three different experimental design criteria: LHS and active learning with the MSE, and the IMSE criteria. To assess the LOO error with the LHS design, we gen-

³² Bayesian active learning can be viewed as a grid-free method that emulates the adaptivity of sparse grids (see, for instance, [Brumm and Scheidegger \(2017\)](#), [Brumm et al. \(2022\)](#)).




Figure 11: Convergence of the GP LOO error for different sizes of the experimental design. We consider LHS, and Bayesian active learning with the MSE and IMSE acquisition functions.

erate a series of experimental designs with incremental sizes $n = \{100, 105, \dots, 200\}$ utilizing LHS. For each size, we compute the error following Eq. (120).

Conversely, for the MSE and IMSE criteria, we begin with an initial design of $n = 100$, sampled via LHS, and progressively add observations by optimizing the acquisition functions. With the MSE criterion, we apply the L-BFGS-B routine to optimize the acquisition function in Eq. (121). On the other hand, for the IMSE criterion, we first randomly sample n_{MC} points from the domain of interest \mathcal{B} to integrate the variance as shown in Eq. (123). We then optimize the Monte Carlo integration using the L-BFGS-B routine. Since both acquisition functions are non-linear and challenging to optimize, we restart from multiple initial conditions, select the best candidate, and add this point to the training data \mathcal{D} for the next iteration.³³

The convergence pattern of each criterion concerning the SCC in 2100 are illustrated in Figure 11. We observe that while the error associated with LHS displays a random behavior, active learning approaches using both MSE and IMSE criteria exhibit a more consistent rate of convergence. Notably, the IMSE criterion demonstrates superior performance compared to the MSE for our specific measure of interest. Consequently, we utilize the IMSE criterion to construct a surrogate model in Section 4. In line with the recommendations of [Le Gratiet et al. \(2017\)](#), a LOO error $\epsilon_{LOO} \leq 0.05$ is considered adequate for calculating the first-order Sobol' indices. However, given our intention to compute the Shapley values as well, we aim for higher accuracy in the surrogate model and thus select $n = 200$ for all subsequent global sensitivity analyses.

³³In our numerical applications, we select $n_{MC} = 128$ and restart 20 times with different initial conditions, which are randomly chosen from the domain of interest \mathcal{B} . In every active learning step, we add the best point from these 20 candidate points to the training data set.

References

- ACKERMAN, F., E. A. STANTON, AND R. BUENO (2010): “Fat Tails, Exponents, Extreme Uncertainty: Simulating Catastrophe in DICE,” *Ecological Economics*, 69, 1657–1665. [4, 5]
- ANDERSON, B., E. BORGONOVO, M. GALEOTTI, AND R. ROSON (2014): “Uncertainty in Climate Change Modeling: Can Global Sensitivity Analysis Be of Help?” *Risk Analysis*, 34, 271–293. [2, 6]
- ANTHOFF, D., R. S. J. TOL, AND G. W. YOHE (2009): “Risk Aversion, Time Preference, and the Social Cost of Carbon,” *Environmental Research Letters*, 4, 024002. [4]
- AZINOVIC, M., L. GAEGAUF, AND S. SCHEIDEGGER (2022): “Deep Equilibrium Nets,” *International Economic Review*, n/a. [2, 4, 5, 13, 40]
- AZINOVIC, M. AND J. ŽEMLIČKA (2023): “Economics-Inspired Neural Networks with Stabilizing Homotopies,” *arXiv preprint arXiv:2303.14802*. [5]
- BARNETT, M. (2023): “Climate change and uncertainty: An asset pricing perspective,” *Management Science*. [4]
- BARNETT, M., W. BROCK, AND L. P. HANSEN (2020): “Pricing Uncertainty Induced by Climate Change,” *The Review of Financial Studies*, 33, 1024–1066. [5]
- BARNETT, M., W. BROCK, L. P. HANSEN, R. HU, AND J. HUANG (2023): “A Deep Learning Analysis of Climate Change, Innovation, and Uncertainty,” *arXiv preprint arXiv:2310.13200*. [5]
- BASU, S. AND B. BUNDICK (2017): “Uncertainty Shocks in a Model of Effective Demand,” *Econometrica*, 85, 937–958. [23]
- BELLMAN, R. E. (1961): *Adaptive Control Processes: A Guided Tour*, Princeton University Press. [3]
- BERGSTRA, J. S., R. BARDENET, Y. BENGIO, AND B. KÉGL (2011): “Algorithms for Hyperparameter Optimization,” in *Advances in neural information processing systems*, 2546–2554. [14]
- BINOIS, M., J. HUANG, R. B. GRAMACY, AND M. LUDKOVSKI (2019): “Replication or Exploration? Sequential Design for Stochastic Simulation Experiments,” *Technometrics*, 61, 7–23. [46]
- BLATMAN, G. AND B. SUDRET (2010): “An Adaptive Algorithm to Build up Sparse Polynomial Chaos Expansions for Stochastic Finite Element Analysis,” *Probabilistic Engineering Mechanics*, 25, 183–197. [44]
- BRETSCHER, L., J. FERNANDEZ-VILLAYERDE, AND S. SCHEIDEGGER (2022): “Ricardian Business Cycle,” *Available at SSRN 4278274*. [13]

- BRUMM, J., C. KRAUSE, A. SCHAAB, AND S. SCHEIDEGGER (2022): “Sparse Grids for Dynamic Economic Models,” . [46]
- BRUMM, J. AND S. SCHEIDEGGER (2017): “Using Adaptive Sparse Grids to Solve High-Dimensional Dynamic Models,” *Econometrica*, 85, 1575–1612. [1, 46]
- BUTLER, M. P., P. M. REED, K. FISHER-VANDEN, K. KELLER, AND T. WAGENER (2014): “Identifying Parametric Controls and Dependencies in Integrated Assessment Models Using Global Sensitivity Analysis,” *Environmental Modelling & Software*, 59, 10–29. [2, 6, 45]
- CAI, Y., K. L. JUDD, T. M. LENTON, T. S. LONTZEK, AND D. NARITA (2015): “Environmental Tipping Points Significantly Affect the Cost-Benefit Assessment of Climate Policies,” *Proceedings of the National Academy of Sciences of the United States of America*, 112, 4606–4611. [4]
- CAI, Y., T. M. LENTON, AND T. S. LONTZEK (2016): “Risk of Multiple Interacting Tipping Points Should Encourage Rapid CO₂ Emission Reduction,” *Nature Climate Change*, 6, 520–525. [4]
- CAI, Y. AND T. S. LONTZEK (2019): “The Social Cost of Carbon with Economic and Climate Risks,” *Journal of Political Economy*, 127, 2684–2734. [1, 2, 3, 4, 5, 7, 8, 17, 24, 25, 31]
- CHEN, H., A. DIDISHEIM, AND S. SCHEIDEGGER (2021): “Deep Structural Estimation : With an Application to Option Pricing,” Tech. rep. [2, 6]
- CHEN, X. AND Q. ZHOU (2014): “Sequential Experimental Designs for Stochastic Kriging,” in *Proceedings of the Winter Simulation Conference 2014*, 3821–3832. [45]
- CROST, B. AND C. P. TRAEGER (2013): “Optimal Climate Policy: Uncertainty versus Monte Carlo,” *Economics Letters*, 120, 552–558. [4, 5, 7]
- (2014): “Optimal CO₂ Mitigation under Damage Risk Valuation,” *Nature Climate Change*, 4, 631–636. [4]
- DE GROOT, O., A. W. RICHTER, AND N. A. THROCKMORTON (2018): “Uncertainty shocks in a model of effective demand: Comment,” *Econometrica*, 86, 1513–1526. [23]
- DEGROOT, M. H. (1970): *Optimal Statistical Decisions*, Wiley. [11]
- EBRAHIMI KAHOU, M., J. FERNÁNDEZ-VILLAYERDE, J. PERLA, AND A. SOOD (2021): “Exploiting Symmetry in High-Dimensional Dynamic Programming,” Working Paper 28981, National Bureau of Economic Research. [4]
- EFTekhARI, A. AND S. SCHEIDEGGER (2022): “High-Dimensional Dynamic Stochastic Model Representation,” *SIAM Journal on Scientific Computing*, 44, C210–C236. [18]
- EFTekhARI, A., S. SCHEIDEGGER, AND O. SCHENK (2017): “Parallelized Dimensional Decomposition for Large-Scale Dynamic Stochastic Economic Models,” in *Proceedings of the Platform for Advanced Scientific Computing Conference*, New York, NY, USA: ACM, PASC ’17, 9:1–9:11. [18]

- EPSTEIN, L. G. AND S. E. ZIN (1989): “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 57, 937. [1, 6]
- FERNÁNDEZ-VILLAYERDE, J., S. HURTADO, AND G. NUÑO (2023): “Financial Frictions and the Wealth Distribution,” *Econometrica*, 91, 869–901. [4]
- FITZPATRICK, L. G. AND D. L. KELLY (2017): “Probabilistic Stabilization Targets,” *Journal of the Association of Environmental and Resource Economists*, 4, 611–657. [4, 10, 28]
- FOLINI, D., A. FRIEDL, F. KUBLER, AND S. SCHEIDEGGER (2023): “The climate in climate economics,” Available at SSRN 3885021. [1, 5, 6, 8, 9, 10, 24, 37]
- GILLINGHAM, K., W. NORDHAUS, D. ANTHOFF, G. BLANFORD, V. BOSETTI, P. CHRISTENSEN, H. MCJEON, AND J. REILLY (2018): “Modeling Uncertainty in Integrated Assessment of Climate Change: A Multimodel Comparison,” *Journal of the Association of Environmental and Resource Economists*, 5, 791–826. [4]
- GLOROT, X. AND Y. BENGIO (2010): “Understanding the difficulty of training deep feed-forward neural networks,” in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, ed. by Y. W. Teh and M. Titterton, Chia Laguna Resort, Sardinia, Italy: PMLR, vol. 9 of *Proceedings of Machine Learning Research*, 249–256. [14]
- GOLOSOV, M., J. HASSLER, P. KRUSELL, AND A. TSYVINSKI (2014): “Optimal Taxes on Fossil Fuel in General Equilibrium,” *Econometrica*, 82, 41–88. [4]
- GOODFELLOW, I., Y. BENGIO, AND A. COURVILLE (2016): *Deep Learning*, MIT Press, <http://www.deeplearningbook.org>. [13]
- HALTON, J. H. (1960): “On the Efficiency of Certain Quasi-Random Sequences of Points in Evaluating Multi-Dimensional Integrals,” *Numerische Mathematik*, 2, 84–90. [45]
- HAN, J., Y. YANG, ET AL. (2021): “DeepHAM: A Global Solution Method for Heterogeneous Agent Models with Aggregate Shocks,” Tech. rep., arXiv preprint arXiv:2112.14377. [5]
- HARENBERG, D., S. MARELLI, B. SUDRET, AND V. WINSCHERL (2019): “Uncertainty Quantification and Global Sensitivity Analysis for Economic Models,” *Quantitative Economics*, 10, 1–41. [2, 5, 6, 16, 18, 20, 23, 44, 45]
- HASSLER, J., P. KRUSELL, AND A. A. SMITH (2016): “Chapter 24 - Environmental Macroeconomics,” Elsevier, vol. 2, 1893–2008. [4]
- HORNIK, K., M. STINCHCOMBE, AND H. WHITE (1989): “Multilayer Feedforward Networks are Universal Approximators,” *Neural Networks*, 2, 359 – 366. [13]
- HWANG, I. C., F. REYNÈS, AND R. S. TOL (2017): “The Effect of Learning on Climate Policy under Fat-Tailed Risk,” *Resource and Energy Economics*, 48, 1–18. [4, 10, 25, 28]

- IOOSS, B. AND C. PRIEUR (2019): “Shapley Effects for Sensitivity Analysis with Correlated Inputs: Comparisons with Sobol’ Indices, Numerical Estimation and Applications,” *International Journal for Uncertainty Quantification*, 9. [21]
- JAYNES, E. T. (1957): “Information Theory and Statistical Mechanics. II,” *Physical review*, 108, 171–190. [18]
- (1982): “On The Rationale of Maximum-Entropy Methods,” *Proceedings of the IEEE*, 70, 939 – 952, cited by: 1118. [18]
- JENSEN, S. AND C. TRAEGER (2013): “Optimally climate sensitive policy: A comprehensive evaluation of uncertainty & learning,” *Department of Agricultural and Resource Economics, UC Berkeley*. [4]
- JENSEN, S. AND C. P. TRAEGER (2014): “Optimal Climate Change Mitigation under Long-Term Growth Uncertainty: Stochastic Integrated Assessment and Analytic Findings,” *European Economic Review*, 69, 104–125. [1, 4, 24, 25]
- JU, N. AND J. MIAO (2012): “Ambiguity, Learning, and Asset Returns,” *Econometrica*, 80, 559–591. [23]
- KELLY, D. L. AND C. D. KOLSTAD (1999): “Bayesian Learning, Growth, and Pollution,” *Journal of Economic Dynamics and Control*, 23, 491–518. [4, 10, 29]
- KELLY, D. L. AND Z. TAN (2015): “Learning and Climate Feedbacks: Optimal Climate Insurance and Fat Tails,” *Journal of Environmental Economics and Management*, 72, 98–122. [4, 10, 11, 25, 28, 29, 30]
- KINGMA, D. P. AND J. BA (2014): “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*. [15, 44]
- KNUTTI, R., M. A. A. RUGENSTEIN, AND G. C. HEGERL (2017): “Beyond Equilibrium Climate Sensitivity,” *Nature Geoscience*, 10, 727–736. [1, 4, 9]
- KOTLIKOFF, L., F. KUBLER, A. POLBIN, AND S. SCHEIDEGGER (2021): “Pareto-improving carbon-risk taxation,” *Economic Policy*, 36, 551–589. [2, 6, 8]
- LE GRATIET, L., S. MARELLI, AND B. SUDRET (2017): “Metamodel-Based Sensitivity Analysis: Polynomial Chaos Expansions and Gaussian Processes,” in *Handbook of Uncertainty Quantification*, ed. by R. Ghanem, D. Higdon, and H. Owhadi, Cham: Springer International Publishing, 1289–1325. [44, 47]
- LEACH, A. J. (2007): “The Climate Change Learning Curve,” *Journal of Economic Dynamics and Control*, 31, 1728–1752. [4, 10]
- LEMOINE, D. AND C. TRAEGER (2014): “Watch Your Step: Optimal Policy in a Tipping Climate,” *American Economic Journal: Economic Policy*, 6, 137–166. [4]
- (2016): “Economics of Tipping the Climate Dominoes,” *Nature Climate Change*, 6, 514–519. [4]

- LENTON, T. M., H. HELD, E. KRIEGLER, J. W. HALL, W. LUCHT, S. RAHMSTORF, AND H. J. SCHELLNHUBER (2008): “Tipping elements in the Earth’s climate system,” *Proceedings of the national Academy of Sciences*, 105, 1786–1793. [2, 5, 8]
- LONTZEK, T. S., Y. CAI, K. L. JUDD, AND T. M. LENTON (2015): “Stochastic Integrated Assessment of Climate Tipping Points Indicates the Need for Strict Climate Policy,” *Nature Climate Change*, 5, 441–444. [4]
- LONTZEK, T. S. AND D. NARITA (2011): “Risk-Averse Mitigation Decisions in an Unpredictable Climate System,” *The Scandinavian Journal of Economics*, 113, 937–958. [4]
- LONTZEK, T. S., D. NARITA, AND O. WILMS (2016): “Stochastic Integrated Assessment of Ecosystem Tipping Risk,” *Environmental and Resource Economics*, 65, 573–598. [4]
- MA, X. AND N. ZABARAS (2010): “An Adaptive High-Dimensional Stochastic Model Representation Technique for the Solution of Stochastic Partial Differential Equations,” *J. Comput. Phys.*, 229, 3884–3915. [18]
- MALIAR, L., S. MALIAR, AND P. WINANT (2021): “Deep Learning for Solving Dynamic Economic Models,” *Journal of Monetary Economics*, 122, 76–101. [4, 42]
- MARREL, A., B. IOOSS, B. LAURENT, AND O. ROUSTANT (2009): “Calculations of Sobol Indices for the Gaussian Process Metamodel,” *Reliability Engineering & System Safety*, 94, 742–751. [19]
- MCKAY, M. D., R. J. BECKMAN, AND W. J. CONOVER (1979): “A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code,” *Technometrics*, 21, 239–245. [45]
- MIFTAKHOVA, A. (2021): “Global sensitivity analysis for optimal climate policies: Finding what truly matters,” *Economic Modelling*, 105, 105653. [2, 6]
- MURPHY, K. P. (2012): *Machine Learning: A Probabilistic Perspective*, The MIT Press. [16]
- (2022): *Probabilistic Machine Learning: Advanced Topics*, MIT Press. [16]
- NORDHAUS, W. D. (2008): *A Question of Balance: Weighing the Options on Global Warming Policies*, Yale University Press. [4, 5]
- (2011): “The economics of tail events with an application to climate change,” *Review of Environmental Economics and Policy*. [29]
- (2017): “Revisiting the Social Cost of Carbon,” *Proceedings of the National Academy of Sciences*, 114, 1518 LP – 1523. [1, 2, 5, 6, 7, 8, 17, 24]
- (2018): “Projections and Uncertainties about Climate Change in an Era of Minimal Climate Policies,” *American Economic Journal: Economic Policy*, 10, 333–360. [4]
- OAKLEY, J. E. AND A. O’HAGAN (2004): “Probabilistic Sensitivity Analysis of Complex Models: A Bayesian Approach,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66, 751–769. [19]

- OWEN, A. B. (2014): “Sobol’ Indices and Shapley Value,” *SIAM/ASA Journal on Uncertainty Quantification*, 2, 245–251. [2, 5, 20]
- PINDYCK, R. S. (2011): “Fat tails, thin tails, and climate change policy,” *Review of Environmental Economics and Policy*. [29]
- (2013): “Climate Change Policy: What Do the Models Tell Us?” *Journal of Economic Literature*, 51, 860–872. [1]
- RASMUSSEN, C. E. AND C. K. I. WILLIAMS (2005): *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*, The MIT Press. [15, 16]
- RENNER, P. AND S. SCHEIDEGGER (2018): “Machine learning for dynamic incentive problems,” *Available at SSRN 3462011*, working paper. [3, 5, 16, 17, 45]
- ROE, G. H. AND M. B. BAKER (2007): “Why Is Climate Sensitivity So Unpredictable?” *Science*, 318, 629–632. [1, 9, 10, 24, 29]
- SACKS, J., S. B. SCHILLER, AND W. J. WELCH (1989): “Designs for Computer Experiments,” *Technometrics*, 31, 41–47. [46]
- SALTELLI, A. AND B. D’HOMBRES (2010): “Sensitivity Analysis Didn’t Help. A Practitioner’s Critique of the Stern Review,” *Global Environmental Change*, 20, 298–302. [5]
- SALTELLI, A., M. RATTO, T. ANDRES, F. CAMPOLONGO, J. CARIBONI, D. GATELLI, M. SAISANA, AND S. TARANTOLA (2007): *Global Sensitivity Analysis. The Primer*, Chichester, UK: John Wiley & Sons, Ltd. [2, 5, 18, 20]
- SANTNER, T. J., B. J. WILLIAMS, AND W. I. NOTZ (2018): *The Design and Analysis of Computer Experiments*, Springer Series in Statistics, New York, NY: Springer. [45]
- SCHEIDEGGER, S. AND I. BILIONIS (2019): “Machine Learning for High-Dimensional Dynamic Stochastic Economies,” *Journal of Computational Science*, 33, 68–82. [6, 13]
- SCHEIDEGGER, S., J. RUST, J. STACHURSKI, F. ISKHAKOV, B. SCHJERNING, Y. YANG, F. KUBLER, J. FERNÁNDEZ-VILLAVERDE, R. A. MILLER, W. NEWEY, S. MISRA, R. SARMIENTO, M. AZINOVIC, H. CHEN, S. MALAMUD, M. N. WHITE, V. SEMENOVA, K. ZHOU, J. PAYNE, J. CHASSOT, G. NUNO, A. PELICAN, H. KASE, AND L. MALIAR (2023): “Econometric Society Summer-school on Deep Learning for Solving and Estimating Dynamic Models,” . [13]
- SHAPLEY, L. S. (1953): “A Value for N-Person Games,” in *Contributions to the Theory of Games (AM-28), Volume II*, ed. by H. W. Kuhn and A. W. Tucker, Princeton: Princeton University Press, 307–318. [20]
- SMITH, R. C. (2014): *Uncertainty Quantification: Theory, Implementation, and Applications*, Society for Industrial and Applied Mathematics. [5]
- SOBOL, I. M. (1967): “On the Distribution of Points in a Cube and the Approximate Evaluation of Integrals,” *USSR Computational Mathematics and Mathematical Physics*, 7, 86–112. [45]

- (2001): “Global Sensitivity Indices for Nonlinear Mathematical Models and Their Monte Carlo Estimates,” *Mathematics and Computers in Simulation*, 55, 271–280. [18, 19]
- SONG, E., B. L. NELSON, AND J. STAUM (2016): “Shapley Effects for Global Sensitivity Analysis: Theory and Computation,” *SIAM/ASA Journal on Uncertainty Quantification*, 4, 1060–1083. [2, 20, 21]
- STERN, N. (2008): “The Economics of Climate Change,” *American Economic Review*, 98, 1–37. [24]
- SUDRET, B. (2008): “Global Sensitivity Analysis Using Polynomial Chaos Expansions,” *Reliability Engineering & System Safety*, 93, 964–979. [18, 19]
- SUN, Y., D. W. APLEY, AND J. STAUM (2011): “Efficient Nested Simulation for Estimating the Variance of a Conditional Expectation,” *Operations Research*, 59, 998–1007. [20]
- TRAEGER, C. P. (2014): “A 4-States DICE: Quantitatively Addressing Uncertainty Effects in Climate Change,” *Environmental and Resource Economics*, 59, 1–37. [4, 17, 37]
- (2021): “ACE – Analytic Climate Economy,” SSRN Scholarly Paper ID 3832722, Social Science Research Network, Rochester, NY. [4]
- VAN DER PLOEG, F. (2021): “Carbon pricing under uncertainty,” *International Tax and Public Finance*, 28, 1122–1142. [4]
- WEBSTER, M., L. JAKOBOVITS, AND J. NORTON (2008): “Learning about Climate Change and Implications for Near-Term Policy,” *Climatic Change*, 89, 67–85. [10]
- WEIL, P. (1989): “The Equity Premium Puzzle and the Risk-Free Rate Puzzle,” *Journal of Monetary Economics*, 24, 401–421. [6]
- WEITZMAN, M. L. (2009): “On modeling and interpreting the economics of catastrophic climate change,” *The review of economics and statistics*, 91, 1–19. [30]
- (2011): “Fat-tailed uncertainty in the economics of catastrophic climate change,” *Review of Environmental Economics and Policy*. [29]
- (2012): “GHG Targets as Insurance Against Catastrophic Climate Damages,” *Journal of Public Economic Theory*, 14, 221–244. [2, 5, 7, 8, 24]
- WINTER, E. (2002): “The Shapley Value,” in *Handbook of Game Theory with Economic Applications*, Elsevier, vol. 3, 2025–2054. [20]
- YOUNES, A., T. A. MARA, N. FAJRAOUI, F. LEHMANN, B. BELFORT, AND H. BEYDOUN (2013): “Use of Global Sensitivity Analysis to Help Assess Unsaturated Soil Hydraulic Parameters,” *Vadose Zone Journal*, 12. [20]
- ZALIAPIN, I. AND M. GHIL (2010): “Another Look at Climate Sensitivity,” *Nonlinear Processes in Geophysics*, 17, 113–122. [9]
- ZHAO, Y., A. BASU, T. S. LONTZEK, AND K. SCHMEDDERS (2023): “The Social Cost of Carbon When We Wish for Full-Path Robustness,” *Management Science*. [5]