

ESSAYS IN COMPUTATIONAL AND CLIMATE ECONOMICS

Joseph Salera Huang

A DISSERTATION

in

Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2024

Supervisor of Dissertation

Jesús Fernández-Villaverde, Howard Marks Presidential Professor of Economics

Graduate Group Chairperson

David Dillenberger, Professor of Economics

Dissertation Committee

Jesús Fernández-Villaverde, Howard Marks Presidential Professor of Economics

Dirk Krueger, Walter H. and Leonore C. Annenberg Professor in the Social Sciences and Professor of Economics

Harold L. Cole, James Joo-Jin Kim Professor of Economics

ESSAYS IN COMPUTATIONAL AND CLIMATE ECONOMICS

COPYRIGHT

2024

Joseph Salera Huang

To my mother, Bishan Huang, for her unwavering love, strength, and wisdom, which have supported and inspired me throughout my journey.

To my father, Bernard J. Salera, whose profound love and mind continue to guide me each day. I only wish you could be here to share this moment with us.

ACKNOWLEDGEMENT

First and foremost, I must acknowledge the support from my parents, Bishan Huang and Bernard J. Salera. My mother, the greatest person on earth, has shown me unconditional love for more than three decades and made me who I am. I admire her intelligence and strength. My father, who unfortunately passed away, blessed me with his loving guidance and intellectual passion that still inspire me today.

Deep gratitude goes to my advisors: Jesús Fernández-Villaverde, Dirk Krueger, and Harold Cole. Their rigorous scholarship and insightful feedback have been instrumental in shaping this dissertation. Their mentorship has taught me not just economics, but how to think critically and conduct meaningful research. I am deeply indebted to them for their generous dedication of time and energy, and for always pushing me to reach higher standards in my research. Beyond their mentorship, Jesús and Dirk's exceptional teaching in the first-year courses laid the foundation of my economic intuition and knowledge, providing me with the tools and frameworks that continue to guide my research.

My academic journey has been significantly shaped by Lars Peter Hansen, Mike Barnett, Amy Boonstra, Fabrice Tourre, Paymon Khorrami, Ruimeng Hu, Buz Brock, and Diana Petrova, whose mentorship, advice, and encouragement have been pivotal throughout. I am particularly indebted to Lars Hansen for taking a chance on me as his research assistant, a defining moment that set me on this academic path, and to Amy Boonstra for her instrumental role in this opportunity and continued support during and after the assistantship. I am also deeply grateful to Paymon and Fabrice, whose guidance during my time as a research assistant helped me develop crucial research skills. The experience of collaborating with Lars, Mike, Ruimeng, and Buz has been transformative, culminating in joint work that forms an essential chapter of this dissertation.

I am particularly grateful to Ornella Darova and Simon Scheidegger for their invaluable insights, detailed feedback, and immense support. I also appreciate the helpful comments and discussions with José Víctor Ríos Rull, Rick Presman, Marlon Azinovic, Alessandro DAVIS, Jordan Peebles, and

Ozgur Seker. I thank the participants at the Macro Student Talk and Macro Lunch groups at the University of Pennsylvania. I acknowledge financial support from the Kleinman Center for Energy Policy.

ABSTRACT

ESSAYS IN COMPUTATIONAL AND CLIMATE ECONOMICS

Joseph Salera Huang

Jesús Fernández-Villaverde

Chapter 1 examines computational methods for solving continuous-time macroeconomic models, comparing grid-based methods enhanced by high-performance computing (HPC) and grid-free approaches using deep learning. While HPC improves grid-based methods' performance, I show how Amdahl's Law constrains their scalability. Deep learning methods, particularly the Deep Galerkin Method, offer an alternative that better handles high-dimensional problems by avoiding the curse of dimensionality. In the subsequent chapters, my co-authors and I provide implementation frameworks for both approaches in solving Hamilton-Jacobi-Bellman equations and demonstrate their applications in macroeconomic and climate economics models.

In chapter 2, using firm-level data, I document five empirical facts that reveal a misalignment between high ESG scores and actual environmental performance. To explain these facts, I develop a novel general equilibrium model with distorted learning. The model quantifies the economic and environmental costs of greenwashing. Key findings include: (1) firms can take up to three times as long to reduce emissions compared to scenarios with perfect information; (2) even after a decade, economies with significant greenwashing exhibit emissions that are 0.01 gigatons (0.1% of total US emissions) higher and are 9% less green than in scenarios with perfect information; and (3) greenwashing can increase social cost of carbon by up to \$18 billion due to delayed emission reductions.

In chapter 3, my co-authors (Michael Barnett, William Brock, Lars Hansen, and Ruimeng Hu) and I develop a neural-network-based computational algorithm that incorporates pseudo-states variables to derive accurate global solutions to high-dimensional, non-linear, dynamic models with jumps and uncertainty aversion. We use this algorithm to study optimal carbon-neutral transition policy

in response to climate change, a setting where endogenous nonlinearities from model uncertainty and jump processes are first-order considerations. Our social planner places substantial value on investment in R&D and green capital, though the significant “social pay-off” of green technological innovation means only R&D investment is significantly augmented by uncertainty concerns, even with uncertainty concerns about jumps in the severity of climate change damages.

TABLE OF CONTENTS

ACKNOWLEDGEMENT	iv
ABSTRACT	vi
LIST OF TABLES	x
LIST OF ILLUSTRATIONS	xi
CHAPTER 1 : APPLICATIONS OF HIGH PERFORMANCE COMPUTING AND MA- CHINE LEARNING IN CONTINUOUS-TIME MACROECONOMICS	1
1.1 Introduction	1
1.2 Grid-based Methods	2
1.3 Grid-free Methods	6
1.4 Conclusion	9
CHAPTER 2 : MACROECONOMIC CONSEQUENCES OF GREENWASHING	10
2.1 Introduction	10
2.2 Literature Review	12
2.3 Empirical Evidence	14
2.4 Model	25
2.5 Calibration	40
2.6 Quantitative Analysis	43
2.7 Conclusion	47
CHAPTER 3 : A DEEP LEARNING ANALYSIS OF CLIMATE CHANGE, INNOVATION, AND UNCERTAINTY	48
3.1 Introduction	48
3.2 A Deep Learning Algorithm for Nonlinear Problems	55
3.3 Climate-Economics Model Example	64

3.4	Model Solution	71
3.5	Numerical Results	72
3.6	Algorithm Validation	84
3.7	Conclusion	84
APPENDIX A : DERIVATIONS OF CHAPTER 2		86
APPENDIX B : DERIVATIONS OF CHAPTER 3		89
BIBLIOGRAPHY		108

LIST OF TABLES

TABLE 2.1	Regression results of (2.1). Δ denotes the change in the variable from the previous year. Emission intensity is defined as scope 1 and 2 emissions divided by revenue. I divide emission intensity into quintiles. The two columns use capital and market capitalization (in this order) to measure size. Capital and market capitalization are in (log) \$ millions. <i>Reported</i> indicates whether the firm reported emissions.	19
TABLE 2.2	Regression results of (2.2). Δ denotes the change in the variable from the previous year. Emission improvement is defined as moving into a lower quintile of emission intensity, which is defined as emissions divided by revenue. In this table, size is measured by capital and market capitalization as in table 2.1. . . .	20
TABLE 2.3	Regression results of (2.3). Δ denotes the change in the variable from the previous year. Emission intensity is defined as scope 1 and 2 emissions divided by revenue. I divide emission intensity into four quantiles. In this table, I use scores to predict emissions one ($t+1$) or two years ($t+2$) ahead. Size is measured by capital or market capitalization as in table 2.1.	22
TABLE 2.4	Regression results of (2.4).	24
TABLE 2.5	Regression results of (2.6). Here, <i>Reported</i> is a dummy variable for whether the firm has reported emissions at least once over the sample period.	27
TABLE 2.6	Calibrated parameters	42
TABLE 2.7	Baseline results: first half contains untargeted and targeted moments; second half contains quantitative implications (change in percents when the economy jumps to the bad state) from the calibration in section 2.5.	43
TABLE 2.8	Model Outcomes after 10 Years	45
TABLE 2.9	Regression results from simulations	46
TABLE B.1	Economic Parameters	103
TABLE B.2	Climate Dynamics and Damages Parameters	103
TABLE B.3	State Variable Initial Values and Ranges	104

LIST OF ILLUSTRATIONS

FIGURE 1.1	Illustration of Amdahl’s Law	5
FIGURE 2.1	Participation in S&P Global’s ESG ratings and environmental score distribution over time	16
FIGURE 2.2	Relationships between environmental scores, firm size, and emissions	26
FIGURE 2.3	Relationships between environmental scores and emission variability	27
FIGURE 2.4	Variables of interest for different model specifications	44
FIGURE 3.1	Implied TCRE coefficients across time and models, capturing carbon and temperature uncertainty. The top figure shows the percentiles for temperature responses to emission pulses for all carbon and temperature model combinations. The bottom figure is a histogram for the exponentially weighted average responses of temperature to an emissions pulse from 144 different models using a rate $\delta = .01$	65
FIGURE 3.2	Range of possible damage functions for two cases with different jump thresholds. The solid lines show the average values, and the shaded regions give the range of possible values for $\exp(-n)$, which measures the proportional reduction of the productive capacity of the economy. The blue line and region show the damage function curvature when the jump occurs at $Y_t = \bar{y} = 2.0$. The red line and region show the damage function curvature when the jump occurs at $Y_t = \underline{y} = 1.5$. The black dashed lines indicate the values of Y_t for the upper and lower jump thresholds for the temperature anomalies.	67
FIGURE 3.3	Intensity function, $r_1 = 1.5$ and $r_2 = 2.5$. With this intensity function, the probability of a jump at an anomaly of 1.6 is approximately .02 per annum, increasing to about .08 per annum at an anomaly of 1.7, increasing further to approximately .18 per annum at an anomaly of 1.8 and then to about one third per annum when the anomaly is 1.9.	74
FIGURE 3.4	Simulated pathways of R&D investment and emissions based on the numerical solutions under different uncertainty penalty configurations. Panel (a) shows the pathway for R&D investment as a fraction of total output and Panel (b) shows the pathway for carbon emissions in gigatons of carbon (GtC). For each panel, results for three cases of model uncertainty are shown: $\xi = \infty$ (red line), $\xi = 0.3$ (green line), and $\xi = 0.1$ (blue line). The trajectories are simulated under the baseline probabilities abstracting from the intrinsic randomness.	76
FIGURE 3.5	Simulated pathways of investment in dirty and green capital stocks based on the numerical solutions under different uncertainty penalty configurations. Panel (a) shows the level of investment in the dirty capital stock and Panel (b) shows the level of investment in the green capital stock. For each panel, results for three cases of model uncertainty are shown: $\xi = \infty$ (red line), $\xi = 0.3$ (green line), and $\xi = 0.1$ (blue line). The trajectories are simulated under the baseline probabilities abstracting from the intrinsic randomness.	78

FIGURE 3.6	Undistorted and distorted climate model distributions. The left plot shows the undistorted and distorted distributions with less misspecification aversion ($\xi = 0.3$). The right plot shows the undistorted and distorted distributions with more misspecification aversion ($\xi = 0.3$). The histograms are calculated at year 40, with the trajectories leading up to year 40 simulated under the baseline probabilities and abstracting from the intrinsic randomness.	80
FIGURE 3.7	Undistorted and distorted damage model distributions. The left plot shows the undistorted and distorted distributions with less misspecification aversion ($\xi = 0.3$). The right plot shows the undistorted and distorted distributions with more misspecification aversion ($\xi = 0.3$). The histograms are calculated at year 40, with the trajectories leading up to year 40 simulated under the baseline probabilities and abstracting from the intrinsic randomness.	80
FIGURE 3.8	Undistorted and distorted technology model distributions. The left plot shows the undistorted and distorted distributions with less misspecification aversion ($\xi = 0.3$). The right plot shows the undistorted and distorted distributions with more misspecification aversion ($\xi = 0.3$). The histograms are calculated at year 40, with the trajectories leading up to year 40 simulated under the baseline probabilities and abstracting from the intrinsic randomness.	81
FIGURE 3.9	Simulated pathways of the cumulative probability of a technology or damage jump occurring based on the numerical solutions under different uncertainty penalty configurations. Panel (a) shows the distorted probability of a technology jump occurring and Panel (b) shows the distorted probability of a damage jump occurring. For each panel, results for three cases of model uncertainty are shown: $\xi = \infty$ (red line), $\xi = 0.3$ (green line), and $\xi = 0.1$ (blue line). The trajectories are simulated under the baseline probabilities abstracting from the intrinsic randomness.	82

CHAPTER 1

APPLICATIONS OF HIGH PERFORMANCE COMPUTING AND MACHINE LEARNING IN CONTINUOUS-TIME MACROECONOMICS

1.1. Introduction

In this chapter, I illustrate the state-of-the-art applications of high performance computing (HPC) and machine learning in continuous-time macro- and climate economics. Computational methods are not new in economics. Judd (1996) and Rust (1996) pioneered the role of computations in macroeconomics. Later, Caldara et al. (2012) and Fernández-Villaverde and Hull (2023) develop novel methods to tackle dynamic programming problems from macroeconomic models with substantial complexity on sophisticated hardware, including quantum annealers. Almost simultaneously, another strand of literature combines continuous-time techniques, often applied in finance, and macroeconomics. The seminal paper by Basak and Cuoco (1998) introduces financial frictions in an economy with heterogeneous agents. Brunnermeier and Sannikov (2014) develop a similar model with occasionally-binding constraints and a bimodal stationary distribution. Eberly and Wang (2012) study a model with heterogeneous capital. Hansen et al. (2024) offer a detailed comparison among this class of models. Its associated software, the computational details of which are documented by Huang et al. (2024), is among the first that solves this class of models on HPC architectures. However, the aforementioned papers remain a grid-based method, where the curse of dimensionality is alleviated to a degree with HPC tools: Han et al. (2018) show that grid-based methods are not computationally feasible when a model has more than three state variables. This limitation motivates a grid-free method. Fernandez-Villaverde et al. (2020), Al-Aradi et al. (2022), and Sirignano and Spiliopoulos (2018) discuss the applications of deep learning as a solution method, where one uses neural networks as a universal function approximator to approximate the solution to a partial differential equation. A detailed discussion of deep learning in economics can be found in Fernández-Villaverde et al. (2024a).

Continuous-time models require the researcher to solve a Hamilton-Jacobi-Bellman (HJB) equation.

Formally, I can state the problem as follows.

Statement of the problem. Given the partial differential equation (PDE) with the following form:

$$0 = K + AV + B \cdot \partial_x V + \text{trace}[CC' \partial_{xx'} V] \quad (1.1)$$

where x is a collection of state variables and the coefficients are nonlinear functions of x , V , and its derivatives:

$$K = K(x, V, \partial_x V, \partial_{xx} V)$$

$$A = A(x, V, \partial_x V, \partial_{xx} V)$$

$$B = B(x, V, \partial_x V, \partial_{xx} V)$$

$$C = C(x, V, \partial_x V, \partial_{xx} V)$$

the researcher aims to find a function V , called a (scaled) value function, that satisfies (1.1) and boundary conditions.

1.2. Grid-based Methods

To solve (1.1), one can use a finite difference approach (Hansen et al. (2024), Huang et al. (2024), Kaplan et al. (2018)), with the following modification: augment the PDE with a false transient (i.e. time derivative):

$$\partial_t V = K + AV + B \cdot \partial_x V + \text{trace}[CC' \partial_{xx'} V]$$

Using finite difference approximation for the time derivative yields:

$$\frac{V^{t+\Delta t} - V^t}{\Delta t} = K + AV + B \cdot \partial_x V + \text{trace}[CC' \partial_{xx'} V] \quad (1.2)$$

To solve this equation, one employs an iterative linearization procedure. The key insight is to

evaluate the nonlinear coefficients using the solution from the previous time step. Specifically,

1. Express the coefficients as functions of the known solution V^t and its derivatives:

$$\begin{aligned}
 K^{(t)} &= K(x, V^t, \partial_x V^t, \partial_{xx} V^t) \\
 A^{(t)} &= A(x, V^t, \partial_x V^t, \partial_{xx} V^t) \\
 B^{(t)} &= B(x, V^t, \partial_x V^t, \partial_{xx} V^t) \\
 C^{(t)} &= C(x, V^t, \partial_x V^t, \partial_{xx} V^t)
 \end{aligned}
 \tag{1.3}$$

2. Replace all instances of V in the right-hand side of equation (1.2) with $V^{t+\Delta t}$, while using the coefficients evaluated at time t :

$$\begin{aligned}
 \frac{V^{t+\Delta t} - V^t}{\Delta t} &= K^{(t)} + A^{(t)} V^{t+\Delta t} + B^{(t)} \cdot \partial_x V^{t+\Delta t} \\
 &\quad + \text{trace}[C^{(t)} C^{(t)'} \partial_{xx'} V^{t+\Delta t}]
 \end{aligned}
 \tag{1.4}$$

This transformation yields a linear PDE in $V^{t+\Delta t}$, which can be solved using standard finite difference methods. The solution algorithm proceeds as follows:

1. Initialize with a guess for $V^{t=0}$
2. Using equations (1.3) and (1.4), solve for $V^{t+\Delta t}$
3. Repeat step 2 until convergence, defined as:

$$\left\| \frac{V^{t+\Delta t} - V^t}{\Delta t} \right\| < \epsilon
 \tag{1.5}$$

where ϵ is a pre-specified tolerance level and $\| \cdot \|$ is an appropriate norm.

This is an implicit finite difference scheme, where the algorithm starts with a terminal condition and walks backward in time. The initial guess $V^{t=0}$ can be an arbitrary smooth function, when the researcher is solving an infinite horizon problem.

An inspection of (1.3) and (1.4) reveals that solving this HJB equation requires one to solve (sparse) linear systems repeatedly, an area where the curse of dimensionality becomes a roadblock. The size of the matrix associated with the linear system scales exponentially with the number of state variables. For example, if one discretizes each state variable with 100 points, one needs to solve a one million by one million matrix each iteration. Furthermore, each iteration produces a new matrix, since $V^{(t)}$ gets updated every time. A naive implementation of LU decomposition has time complexity $\mathcal{O}(n^3)$: using a off-the-shelf program to perform LU decomposition repeatedly is not a viable solution.

To alleviate the curse of dimensionality, the researcher can apply HPC tools so that matrix factorization is carried out over several computing units (i.e. CPUs). For example, PARDISO (Pasadakis et al. (2023), Eftekhari et al. (2021), Gaedke-Merzhäuser et al. (2023)) is a sparse linear solver that employs shared-memory and distributed-memory multiprocessors. Hansen et al. (2018) find that using PARDISO can reduce the computational time of solving the linear system by a factor of almost 10 by using 28 CPUs.

Moreover, in a grid-based method, the researcher should design the model solver so that most, if not all, parts of the solver can be parallelized. This is illustrated by Amdahl's Law (Amdahl, 1967), which quantifies the maximum potential speedup when parallelizing a program. Denote the fraction of a program that can be parallelized as p (where $0 \leq p \leq 1$), and the number of processors as n . The theoretical speedup $S(n)$ is given by:

$$S(n) = \frac{1}{(1-p) + \frac{p}{n}} \tag{1.6}$$

This equation reveals the speed-up from parallelization and the number of cores: $S(n)$ decreases in n , as shown in figure 1.1.

This equation also reveals a fundamental limitation: even with an infinite number of processors ($n \rightarrow \infty$), the maximum achievable speedup is bounded by:

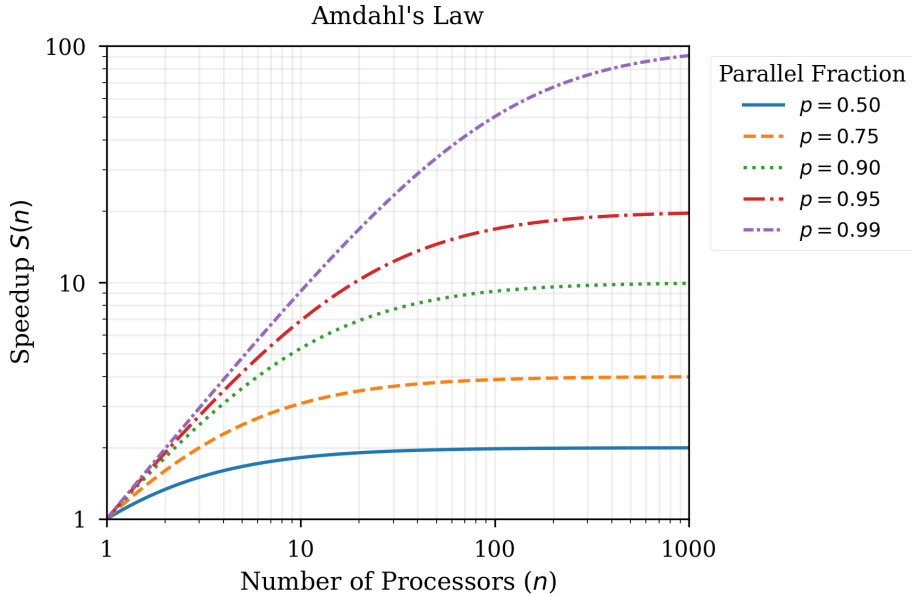


Figure 1.1: Illustration of Amdahl's Law

$$\lim_{n \rightarrow \infty} S(n) = \frac{1}{1-p}$$

For instance, if 90% of a program can be parallelized ($p = 0.9$), the maximum theoretical speedup is 10, regardless of how many additional processors are added. This underscores the critical importance of maximizing the parallelizable portion of the code in PDE solvers, as the non-parallelizable sections fundamentally limit the potential performance gains from parallel computing resources. As seen in figure 1.1, the lines eventually plateau when $p < 1$.

This limitation, together with Amdahl's Law, is one of several factors motivating the adoption of deep learning methods in solving this class of models. Neural networks are particularly well-suited for parallel computation because their core operations – matrix multiplications – are inherently parallelizable. When training neural networks, the parallel fraction p approaches 1, as nearly all computations can be distributed across multiple processing units. Modern GPUs are designed to exploit this parallelism, containing thousands of cores optimized for simultaneous matrix operations.

1.3. Grid-free Methods

Grid-free methods are largely based on neural networks. A neural network is a function approximator that maps inputs $x \in \mathbb{R}^{d_{\text{in}}}$ to outputs $y \in \mathbb{R}^{d_{\text{out}}}$ through a series of nonlinear transformations. The most basic form, a feedforward neural network, consists of layers of neurons where each layer performs an affine transformation followed by a nonlinear activation function. For a network with L layers, each layer l computes:

$$h^{(l)} = \sigma(W^{(l)}h^{(l-1)} + b^{(l)}) \quad (1.7)$$

where $h^{(l)}$ is the output of layer l , $W^{(l)}$ is a weight matrix, $b^{(l)}$ is a bias vector, and $\sigma(\cdot)$ is a nonlinear activation function such as the rectified linear unit (ReLU): $\sigma(x) = \max(0, x)$. The input layer is $h^{(0)} = x$, and the final layer $h^{(L)}$ produces the output y .

The network is trained by minimizing a loss function \mathcal{L} that measures the discrepancy between the network's predictions and target values. Given a dataset of input-output pairs $\{(x_i, y_i)\}_{i=1}^N$, the training objective is:

$$\min_{W^{(l)}, b^{(l)}} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(f_{\theta}(x_i), y_i) \quad (1.8)$$

where f_{θ} represents the neural network with weights $\theta = \{W^{(l)}, b^{(l)}\}_{l=1}^L$. This optimization is typically solved using stochastic gradient descent and backpropagation.

The universal approximation theorem (Cybenko, 1989; Hornik, 1991) provides the theoretical foundation for using neural networks to approximate PDEs. The theorem states that a feedforward neural network with a single hidden layer and sufficient width can approximate any continuous function on a compact domain to arbitrary precision. More formally, for any continuous function $g : \mathbb{R}^{d_{\text{in}}} \rightarrow \mathbb{R}$ and any $\epsilon > 0$, there exists a neural network f_{θ} such that:

$$\sup_{x \in K} \|f_\theta(x) - g(x)\| < \epsilon \quad (1.9)$$

for any compact subset $K \subset \mathbb{R}^{d_{\text{in}}}$. Since solutions to PDEs are continuous functions (under appropriate regularity conditions), neural networks can, in principle, approximate them to arbitrary accuracy. The inherent parallelism in neural networks can be seen in (1.7). Suppose every layer has m neurons and k outputs. $W^{(l)}$ in (1.7) has size $m \times k$ and b has size $k \times 1$. The entire computation $W^{(l)}h^{(l-1)} + b^{(l)}$, requiring $\mathcal{O}(mk)$, can be executed simultaneously on a GPU, effectively achieving a parallel fraction of $p > 0.99$. According to equation 1.6, this enables a theoretical speedup of over a factor of 100. Most of the computational time is spent on passing data between CPUs and GPUs, but with optimization routines in Tensorflow, it doesn't amount to a significant cost.

1.3.1. Approximation criterion

The Deep Galerkin Method (DGM) literature (Fernandez-Villaverde et al., 2020; Al-Arabi et al., 2022; Sirignano and Spiliopoulos, 2018) computes the L^2 norm of (1.1) as the loss. In each training step, the DGM algorithm evaluates the loss function below, which combines the PDE residual with boundary conditions:

$$\mathcal{L}(\theta) = \mathcal{L}_{\text{PDE}}(\theta) + \mathcal{L}_{\text{BC}}(\theta) \quad (1.10)$$

where \mathcal{L}_{PDE} is the PDE residual:

$$\mathcal{L}_{\text{PDE}}(\theta) = \frac{1}{N_{\text{int}}} \sum_{i=1}^{N_{\text{int}}} |K + Af_\theta(x_i) + B \cdot \nabla_x f_\theta(x_i) + \text{trace}[CC' \nabla_{xx'} f_\theta(x_i)]|^2 \quad (1.11)$$

and \mathcal{L}_{BC} enforces the boundary conditions:

$$\mathcal{L}_{\text{BC}}(\theta) = \frac{1}{N_{\text{BC}}} \sum_{j=1}^{N_{\text{BC}}} |f_\theta(x_j^{\text{BC}}) - g(x_j^{\text{BC}})|^2 \quad (1.12)$$

Here, $\{x_i\}_{i=1}^{N_{\text{int}}}$ are sampled points in the interior of the domain, $\{x_j^{\text{BC}}\}_{j=1}^{N_{\text{BC}}}$ are points on the boundary and $g(x)$ specifies the boundary values.

The DGM algorithm seeks to minimize (1.10) to zero. A detailed description of the algorithm is provided in section 3.2.5.

One theoretical limitation of this algorithm is that literature has not offered any guarantee that minimizing (1.10) to zero is equivalent to solving the PDE. That is, let t be the iterant on the training steps and θ^t be the weights at the t th step. I cannot claim, in general, as $t \rightarrow \infty$

$$\mathcal{L}(\theta^t) \rightarrow 0 \quad \text{and} \quad f_{\theta^t} \rightarrow V \tag{1.13}$$

where V is the solution of the PDE. However, nevertheless, under strong conditions (e.g. Lipschitz conditions on the coefficients of the PDE), Sirignano and Spiliopoulos (2018) show that (1.13) holds for a class of quasilinear parabolic PDEs. Empirical tests (Son et al., 2021) of the DGM algorithm for PDEs with known analytical solutions suggest that the applications of DGM are broader than those allowed by the strong conditions imposed in Sirignano and Spiliopoulos (2018).

1.3.2. Scalability

Neural networks alleviate the curse of dimensionality since they are scalable with respect to the number of inputs. Unlike grid-based methods such as finite difference, which stores the value of V at each point, neural networks learn an approximation of V through its weights. During training and inference, the researcher only needs to compute matrix operations on a GPU over, for example, a few hundred points. The universal approximation theorem and empirical tests of neural networks show that neural networks can approximate almost any function without a very sophisticated architecture. Therefore, expanding the number of state variables (i.e. the number of inputs) will not significantly increase computational costs. This property allows the researcher to explore high-dimensional models. Furthermore, as my co-authors and I do in chapter 3, the researcher can add parameter values as inputs to neural networks. Training a neural network is the same as solving a model over a huge set of parameterizations. This "trick" substantially reduces

the computational time needed for calibration or estimation. I leave the detailed discussion of this “trick” in section 3.2.5.

One drawback of deep learning is the lack of formal theory in determining the hyperparameters, such as learning rates, activation functions, etc. Often times the researcher has to choose these hyperparameters through trial and error. However, in practice, my co-authors and I find that we can find a satisfactory solution with standard or default hyperparameters used in the literature. We provide our implementation details in section B.5.

1.4. Conclusion

I explore two paradigms of solving HJB equations that arise from continuous-time models. Grid-based methods benefit substantially from HPC architectures, especially in the face of curse of dimensionality. However, Amdahl’s Law imposes a limitation to this paradigm. Therefore, to handle more complex models, the researcher may need to use grid-free methods such as deep learning. In chapter 2, I develop a model that can be solved in a grid-based method. In chapter 3, my co-authors and I apply deep learning to solve a substantially more complicated model.

CHAPTER 2

MACROECONOMIC CONSEQUENCES OF GREENWASHING

2.1. Introduction

Financial markets and public institutions reward firms that claim to be green. Regarding financial markets, Hartzmark and Sussman (2019) provide causal evidence that for mutual funds, being labeled as low sustainability leads to outflows of more than \$12 billion whereas those deemed highly sustainable attract inflows over \$24 billion. Additionally, Pastor et al. (2022) find the existence of “ESG flows” and a -1% “greenium”, suggesting that firms with high ESG scores have a higher valuation. On the public side, pension funds in New York City and California have announced net-zero emissions goals for their portfolios and select investments based on companies’ alignment with their climate goals. One explanation for this green preference is that green firms may serve as hedges against transition risk, an example of which is a carbon tax (Giglio et al., 2020): when the advent of a carbon tax coincides with reduced production, firms that are less reliant on brown capital (e.g. fossil fuels) pay a higher cash flow and serve as a hedge. Before the arrival of a carbon tax, green firms then have a lower cost of capital, because investors are willing to pay a premium for the inherent hedge. Literature (Fried et al., 2021) has examined the impact of this green preference on the transition to carbon net-zero, since this green preference incentivizes firms to disinvest from brown capital.

Greenness, nevertheless, cannot be directly observed. To assess greenness, investors increasingly rely on ESG rating agencies, which provide a composite score based on a firm’s environmental, social, and governance practices. Literature has found that these scores are significant in explaining firms’ stock returns (Engle et al., 2020). Reliance on these scores, at the same time, opens the door for greenwashing, a deceptive practice where firms falsely present themselves as environmentally friendly without actually reducing emissions. In this paper, I empirically assess the effectiveness of environmental scores and the macroeconomic consequences of greenwashing. I develop a model incorporating *distorted learning* to capture greenwashing. After that, I quantify the economic costs

of greenwashing by computing the increased social cost of carbon that arises from delayed emission reductions.

First, I show five facts regarding environmental scores and environmental performance, measured by emissions in section 2.3 and figures 2.2 and 2.3: (1) firms that are large, decline in size, and report emissions tend to have higher scores; (2) scores are not backward looking: improvement in emissions from the past don't predict higher scores; (3) scores are not forward looking: scores don't predict future emission reductions; (4) higher environmental scores are associated with increased incidence of misleading communication, emissions-related issues, and broader environmental concerns; and (5) the documented emissions of firms that receive high ESG scores tend to exhibit more variance.

Motivated by these facts, I develop a model with distorted learning to quantify the impact of greenwashing, which, as I later quantitatively show, delays emission reductions and increases the social cost of carbon. The mechanism in my model is as follows. My economy starts from a good climate state, with some possibility of jumping to a bad climate state, where production is damaged (hence less consumption) and firms need to pay a carbon tax based on their amount of brown capital, which I define as capital that produces emissions. Firms with less brown capital (i.e. greener firms) pay less in taxes and are hedges against climate risk. This serves as the basis for a *greennium*: households demand a lower required rate of return from greener firms. A key financial friction arises because households can only observe firms' aggregate capital, not their capital composition. Households, then, need to infer firms' greenness from reported emissions, through a Kalman-Bucy filter. I model greenwashing as a way to misspecify the noise of reported emissions, which, in continuous time through the Girsanov theorem, is the same as distorting the drift of the filter (Hansen and Sargent, 2011).

Consistent with empirical facts, my model predicts that large firms greenwash more and appear greener to households. Second, green appearances don't predict reductions in emissions. Furthermore, the model allows me to quantify the key economic quantities along transition paths, as firms become greener over time. In particular, I simulate the economy under different scenarios (i.e. with and without perfect observations of greenness and varying degrees of greenwashing), and analyze

how firms reallocate to green capital and reduce emissions. I find that when households can perfectly observe firms' greenness, firms quickly remove its brown capital and reduce emissions, whereas when firms can greenwash, firms can take three times as long to achieve the same emission reductions. Moreover, when greenwashing is cheap, even after 10 years, firms can be 0.01 gigatons (0.1% of total US emissions) behind in emission reductions (compared to scenarios with mild and no greewashing) and 9% browner than the scenario where households can observe greenness. Laslty, greenwashing can increase the social cost of carbon by up to \$18 billion, as firms delay their emissions reductions.

This paper makes several key contributions to the literature. First, it provides a novel theoretical framework that integrates both learning and model distortion into a general equilibrium model. Second, it offers a quantitative assessment of the economic and environmental costs of greenwashing, which has been largely absent from previous studies. Third, it sheds light on the interplay among green investing, environmental performance, and emissions reporting, contributing to the ongoing debate about the effectiveness of voluntary corporate environmental disclosures.

The rest of the paper is organized as follows. Section 2.2 reviews the related literature. Section 2.3 presents the empirical facts that motivate my model. Section 2.4 details the theoretical model and its equilibrium properties. Section 2.5 describes the calibration. Section 2.6 discusses quantitative implications. Section 2.7 concludes.

2.2. Literature Review

There is a plethora of empirical evidence supporting the existence of greenwashing and the connection between size and green appearance. For example, Boffo and Patalano (2020) from OECD find an almost non-existent correlation between the E(nvironmental) score and carbon intensity. Duchin et al. (2022) provide evidence that firms' pollution levels do not decline although they highlight their divesture efforts in conference calls. Cohen et al. (2023) find that firms with low ESG scores tend to innovate more in green patents. Sastry et al. (2024) find that, using administrative data on European banks, net-zero commitments do not lead to divestments from carbon-intensive projects. Similarly, Elmalt et al. (2021) argue that ESG scores often reflect firms' claims rather than their actual impact on reducing carbon emissions. Moreover, Acharya et al. (2024) find that

large firms are more likely to make green pledges. Lanteri and Rampini (2023) discuss the role of financial constraints in adopting clean technologies and provide evidence that larger firms operate with higher energy efficiency because they can afford higher down payments. Literature, nevertheless, does not offer a general equilibrium model that involves greenwashing. Schmittmann and Gao (2023) model information asymmetry in green bond issuance, but it is not a general equilibrium model. Nor is it quantitative. My paper is the first to offer a quantitative, general equilibrium model that incorporates greenwashing, to the best of my knowledge.

The two-capital model structure in continuous time is first developed by Wang and Eberly (2010). Barnett et al. (2020) and Barnett et al. (2023a) incorporate a climate sector into the two-capital framework. My model is very close to Barnett et al. (2023a). As discussed in section 2.4, the only aggregate risk is the climate transition risk, which I model as a Poisson jump to a state with lowered production and a carbon tax. Using Poisson jumps as climate risk or damages has been done by Barnett et al. (2023a), Hong et al. (2023a), and Hong et al. (2023b). Fried et al. (2021) and Carattini et al. (2023) develop quantitative models for the effects of a carbon tax. Literature has studied the effects of learning in continuous-time models. For example, Pástor and Veronesi (2012) model the learning of government policy. Hong et al. (2023a) provide a model where agents learn about the arrival rate of a climate disaster. The derivations of the Kalman-Bucy filter that I use are well-described in Liptser and Shiryaev (2001) and Øksendal (1998). Nevertheless, these papers do not incorporate distortion into learning. Hansen and Sargent (2011) discuss misspecified dynamics. Specifically, they propose a min-max framework in which the decision maker avoids the worst possible outcome. The set of possible outcomes is defined as the set of admissible probability measures, over which the nature chooses to minimize the decision maker’s welfare. By applying the Girsanov theorem, probability measures are isomorphic to drift distortions (with some loss of generality). I adopt this framework to my model, where firms distort the drift of the noise in their reported emissions, which is, at the same time, a signal that households use to infer firms’ greenness. Imperfect learning, where one forms estimates that differ those implied by Bayes’ rule, has also been studied in literature. For example, Maxted (2023) embeds “diagnostic expectations” in a macroeconomic model. Gabaix (2019) summarizes ways through which households can have

“bounded rationality”.

The equilibrium concept in my paper is a mean-field game (MFG) equilibrium, discussed in Lasry and Lions (2007). Achdou et al. (2022) apply this concept to study the wealth distribution. In my model, I have a distribution of firms, whose controls depend on moments of the economy, which, in turn, depend on other firms’ decisions. In section 2.4.4, I formally define the equilibrium.

2.3. Empirical Evidence

I present five facts regarding the discrepancy between high environmental scores and environmental performance. These facts will guide the development of my model and serve as a reference point for its evaluation.

Data. I use ESG scores and emissions data from S&P Global, a rating agency that started to give out ESG scores in 2013 and provides emissions data through its Trucost division. S&P Global’s influence can be seen in its widespread adoption by major financial institutions, including BlackRock and Vanguard, the two largest asset managers globally. It is also well-studied by academics such as Berg et al. (2019), Bolton and Kacperczyk (2021) and Hartzmark and Shue (2022).

An ESG score, which ranges from 0 to 100, is a composite metric that evaluates a firm’s performance across its environmental, social and governance dimensions. Firms’ ESG credentials are assessed through S&P Global’s annual Corporate Sustainability Assessment (CSA), wherein a selected group of invited companies provide comprehensive disclosures on their operations and practices. Questions range from emission targets, measures taken to achieve these targets, CEO’s compensation, etc. Question-level scores are aggregated into criterion-level scores for the most material ESG themes per industry, which further roll up into separate Environmental, Social and Governance dimension scores. In this paper, I focus on the Environmental (E) dimension. Going forward, unless otherwise specified, *scores* refer to the Environmental scores.

Trucost’s firm-level emissions data includes both direct (scope 1) and indirect (scopes 2 and 3) emissions. Scope 1 emissions are those emitted directly from sources controlled or owned by the company, while Scope 2 emissions cover indirect emissions from the generation of purchased elec-

tricity consumed by the company. Scope 3 emissions include all other indirect emissions that occur in a company's value chain, such as those associated with the production and transportation of purchased goods and services. I will focus on scopes 1 and 2 as in Hartzmark and Shue (2022), since they are the scopes that firms have control over. The Trucost data is compiled through a combination of self-reported information and proprietary estimation models. In addition, I will use emission intensity, which is emissions scaled by total revenue, to adjust for firm size, since larger firms may emit more due to solely their scale. Trucost imputes firms' emissions if firms do not report. Going forward, I will refer to the emissions data in Trucost as "documented emissions", since in section 2.4 I will use "reported emissions" to refer to the emissions firms report.


Data on firms' fundamentals is obtained from Compustat, as is standard in the literature. The frequency of analysis is annual and the coverage only includes firms in the U.S. Lastly, I use data from RepRisk to characterize negative events (such as bad press) associated with greenwashing. RepRisk is a leading data science company that systematically captures and analyzes risk incidents related to environmental, social, and governance (ESG) issues from public sources and news media (Twitter, Facebook, press releases, earnings calls, news channels, etc.). Categories of risk incidents include misleading communication and emissions. Glossner (2017) uses RepRisk data to measure ESG risk and finds that statistically significant impact on the returns of firms with high ESG risk, suggesting that RepRisk captures material negative events.

In figure 2.1, I show the participation in S&P Global's ESG ratings and the distribution of environmental scores over time. Over the decade, more firms are participating in ESG ratings, going from less than 300 in 2014 to over 900 in 2023. The distribution of scores is very skewed, with a long right tail. The skewed distribution is persistent over time.

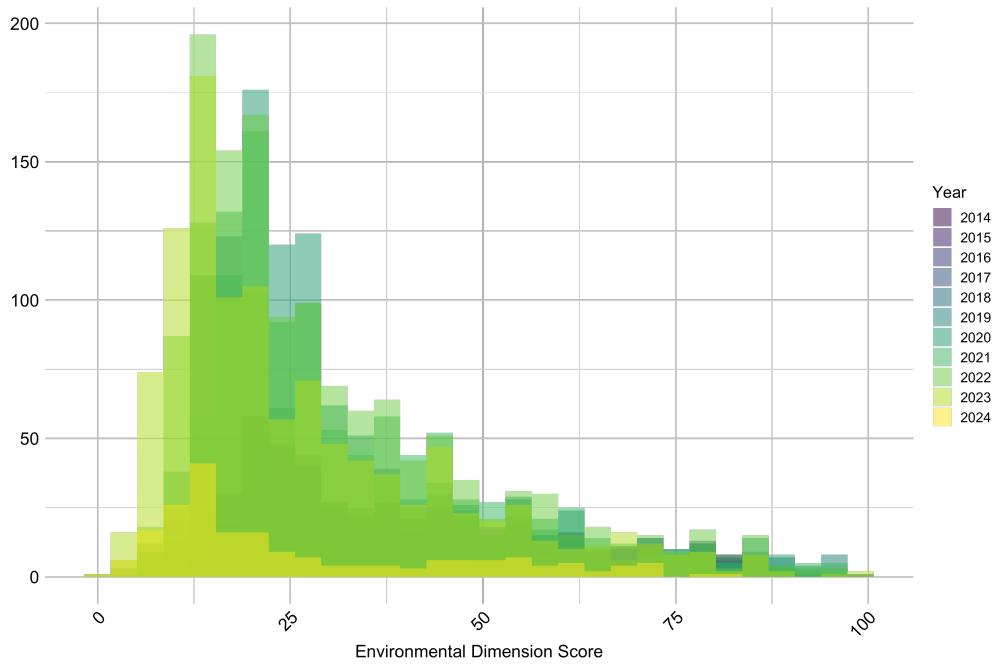
Below, I illustrate the five facts, My sample covers all the U.S. publicly traded firms that are given an ESG score by S&P Global from 2014 to 2024 across all industries.

Fact 1: firms that receive high scores tend to be large, declining, and report emissions.

Firm size, measured by capital or market equity, is a strong predictor of scores. The effect is



(a) Number of Firms in Sample by Year



(b) Distribution of Environmental Dimension Scores Over Time

Figure 2.1: Participation in S&P Global's ESG ratings and environmental score distribution over time

statistically significant and robust to various controls. To show this, I regress scores on size and, in addition, change in size, the quintiles of emission intensity and whether the firm improves its emission intensity from the previous year or two years ago, all at the firm level. Quintiles are determined by fixing firms' industry and reported year. Specifically, I estimate the following regression:

$$\begin{aligned}
\text{score}_{it} = & \beta_1 \text{size}_{it} + \beta_2 \Delta \text{size}_{it} + \beta_3 \text{reported}_{it} + \beta_4 \text{emission quintile}_{it} \\
& + \beta_5 \text{emission quintile}_{it} \times \text{reported}_{it} + \beta_6 \text{improved emissions (t - 1)}_{it} \\
& + \beta_7 \text{improved emissions (t - 2)}_{it} + \gamma_{\text{industry FE}_i} + \gamma_{\text{time FE}_t} + \epsilon_{it}
\end{aligned} \tag{2.1}$$

where i denotes each firm and t each year. Size_{it} is either capital or market capitalization (in logs). I use two different measures of size for robustness. Industry fixed effects are included partly to control for the relative importance of the environmental score in the composite ESG score. Reported_{it} is a dummy variable on whether the firm has reported emissions at least once in the past. If emissions are not reported, Trucost uses its own proprietary estimation models (i.e. they impute the firm's emissions). *Improved emissions (t-1)* (and *t-2*) is defined as whether the firm moves down a quintile in emission intensity from the previous year (and two years ago).

The regression results are summarized in table 2.1 for both measures of size. A few remarks are in order. First, size is a strong predictor of the environmental score, even after controlling for fixed effects and the aforementioned regressors. The statistical significance is robust to both measures of size. Another robust and significant predictor of scores is the change in size. Firms that decline in size tend to have higher scores. Lastly, if the firm reports its emissions (such as responding to the CSA), it has a higher chance of receiving a higher score.

Moreover, the effect of firms' current emission intensity (i.e. the coefficients on quintiles) on their scores are ambiguous and nonrobust (coefficients are non-monotonic and some of them are not significant). In summary, regression results in tables 2.1 provide evidence for fact 1: firms that are large, have a decline in size, and report emissions tend to have higher scores. This fact is visualized in figure 2.2. The relationship between size and scores is striking with almost non-existent standard

errors. One can also see a negative relationship between change in size and scores.

Fact 2: improving emissions from the past year or two years ago does not predict higher scores. Table 2.1 also shows that scores are not backward-looking. Moving down a quintile in emission intensity from the previous year or two years ago does not predict a higher score. Put differently, coefficients on *Improved emissions (t-1)* (and *t-2*) are not statistically significant.

Fact 3: scores don't predict future emissions. In this fact, I show that scores are not forward-looking either: they don't predict future emissions. I estimate the following probit model:

$$P(\text{improved emissions}_{i,t+k} = 1) = \Phi(\beta_0 + \beta_1 \text{score}_{it} + \beta_2 \Delta \text{score}_{it} + \beta_3 \text{size}_{it} + \beta_4 \Delta \text{size}_{it} + \beta_4 \gamma_{\text{industry FE}_i} + \beta_5 \gamma_{\text{time FE}_t}) \quad (2.2)$$

I include size and change in size as predictors to prevent omitted variable bias. They strongly correlate with scores (facts 1 and 2) and, *a priori*, one may reasonably argue that they influence emissions. I estimate two sets of regressions, one to predict change in emissions in one year and the other to predict change in emissions in two years ($k = 1$ and $k = 2$).

Results in table 2.2 indicate that environmental scores and change in environmental scores are not significant predictors of future emissions. Furthermore, despite explaining more than 30% of the variance in scores, size does not predict future emissions (the significance of market capitalization is most likely an artifact, since it is the only size measure that has significance in the four models). Change in size does not explain future emissions either. Moreover, those firms that report emissions are less likely to improve their emissions two years ahead. This is rather surprising, since reporting firms are more likely to receive high scores and receive more attention.

Since the environmental score is a weighted average of pillar scores that assess a firm's environmental credentials, I also use one of the pillar scores, the *emission intensity scopes 1 and 2 scores*, to predict emissions. These pillar scores aim to assess firms' emission intensity and the pertinent section in

	Environmental score	
	(1)	(2)
Size	5.90*** (0.39)	5.17*** (0.43)
Δ Size	-7.25*** (0.95)	-4.90*** (0.62)
Emission Intensity:		
Quintile 2	-2.90** (1.27)	-3.14** (1.26)
Quintile 3	-1.31 (0.94)	-1.32 (1.14)
Quintile 4	-1.45 (0.94)	-1.29 (1.23)
Quintile 5	-0.39 (1.06)	0.73 (1.14)
Reported	14.28*** (1.84)	14.84*** (1.96)
Interactions:		
Q2 \times Reported	2.18 (2.72)	3.09 (2.97)
Q3 \times Reported	-1.52 (2.14)	-0.39 (2.25)
Q4 \times Reported	-2.31 (2.38)	-1.16 (2.47)
Q5 \times Reported	0.48 (2.54)	1.33 (2.49)
Improved emissions:		
Improved emissions ($t - 1$)	-0.17 (0.57)	-0.10 (0.61)
Improved emissions ($t - 2$)	-0.38 (0.85)	-0.24 (0.86)
Time and Industry FE	Yes	Yes
Observations	7,376	7,373
R ²	0.433	0.409

Note: *p<0.1; **p<0.05; ***p<0.01

Table 2.1: Regression results of (2.1). Δ denotes the change in the variable from the previous year. Emission intensity is defined as scope 1 and 2 emissions divided by revenue. I divide emission intensity into quintiles. The two columns use capital and market capitalization (in this order) to measure size. Capital and market capitalization are in (log) \$ millions. *Reported* indicates whether the firm reported emissions.

	Probability of Emissions Improvement			
	$t + 1$	$t + 2$	$t + 1$	$t + 2$
	(1)	(2)	(3)	(4)
Env. score	-0.003* (0.002)	-0.002 (0.002)	-0.003* (0.002)	-0.002 (0.002)
Δ Env. score	-0.001 (0.003)	-0.004 (0.003)	-0.001 (0.003)	-0.004 (0.003)
Size	0.017 (0.019)	0.032 (0.020)	0.014 (0.018)	0.048** (0.020)
Δ Size	0.143 (0.130)	0.123 (0.133)	-0.021 (0.059)	-0.053 (0.068)
Reported	-0.085 (0.058)	-0.149** (0.059)	-0.084 (0.058)	-0.157*** (0.059)
Time and Industry FE	Yes	Yes	Yes	Yes
BIC	4892.29	4523.49	4893.62	4520.32
Observations	7,376	7,376	7,373	7,373

Note: *p<0.1; **p<0.05; ***p<0.01

Table 2.2: Regression results of (2.2). Δ denotes the change in the variable from the previous year. Emission improvement is defined as moving into a lower quintile of emission intensity, which is defined as emissions divided by revenue. In this table, size is measured by capital and market capitalization as in table 2.1.

the CSA asks that firms report emissions¹. I estimate the following Probit model:

$$\begin{aligned}
 P(\text{improved emissions}_{i,t+k} = 1) = & \Phi(\beta_0 + \beta_1 \text{emission intensity pillar score}_{it} \\
 & + \beta_2 \Delta \text{emission intensity pillar score}_{it} + \beta_3 \text{size}_{it} \quad (2.3) \\
 & + \beta_4 \Delta \text{size}_{it} + \beta_4 \gamma_{\text{industry FE}_i} + \beta_5 \gamma_{\text{time FE}_t})
 \end{aligned}$$

where emission intensity pillar score_{it} is the average of the emission intensity scopes 1 and 2 scores for firm *i* in year *t*. The only difference between this regression and the previous one (2.2) is that I use emission intensity pillar scores instead of the Environmental score.

Regression results in table 2.3 show qualitatively similar results as in table 2.2: even emission intensity scores, which are designed to evaluate emission intensity, don't predict future emission reductions. If anything, a high emission intensity score predicts a lower chance of emission reduction. The non-existent relationship between scores and emissions is seen in figure 2.2.

Fact 4: higher environmental scores are associated with increased incidence of misleading communication, emissions-related issues, and broader environmental concerns. To measure misleading communication, emissions, and environmental issues, I use the RepRisk dataset. For each firm in a given year, I calculate the number of RepRisk events that belong to these categories: misleading communication, GHG emissions, and environmental issues overall, where the first two are subcategories of the last. I investigate whether firms with high scores will lead to more RepRisk events in the following year. I then estimate a negative binomial regression (since the dependent variable consists of count data):

$$\begin{aligned}
 \mathbb{E} [\# \text{ RepRisk events}_{i,t+1}] = & \exp \left[\beta_1 \text{market cap}_{it} + \beta_2 \text{score}_{it} + \beta_3 \text{reporting score}_{it} \right. \\
 & \left. + \beta_4 \text{emission quintile}_{it} + \gamma_{\text{industry FE}_i} + \gamma_{\text{time FE}_t} + \varepsilon_{i,t} \right] \quad (2.4)
 \end{aligned}$$

¹The CSA Handbook released by S&P Global states that the question rationale as “Producing more with less material is essential for many industries affected by the increasing scarcity of natural resources. Operational Eco-Efficiency can enhance companies’ competitiveness through reduced costs and environmental liabilities. It can also mean companies are better prepared for future environmental regulations. The key focus is on the inputs and outputs of business operations, and the assessment of trends in the consumption of natural resources and the production of environmental waste products specific to each industry.”

	Probability of Emissions Improvement			
	$t + 1$	$t + 2$	$t + 1$	$t + 2$
	(1)	(2)	(3)	(4)
Emission intensity pillar score	-0.006** (0.003)	-0.007** (0.003)	-0.006** (0.003)	-0.008** (0.003)
Δ Emission intensity pillar score	0.001 (0.004)	-0.004 (0.004)	0.001 (0.004)	-0.003 (0.004)
Size	0.015 (0.018)	0.037* (0.020)	0.012 (0.018)	0.048** (0.020)
Δ Size	0.161 (0.132)	0.142 (0.135)	-0.016 (0.060)	-0.058 (0.069)
Reported	-0.088 (0.059)	-0.121** (0.059)	-0.087 (0.058)	-0.129** (0.059)
Time and Industry FE	Yes	Yes	Yes	Yes
BIC	4727.33	4349.77	4729	4347.71
Observations	7,113	7,113	7,110	7,110

Note:

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 2.3: Regression results of (2.3). Δ denotes the change in the variable from the previous year. Emission intensity is defined as scope 1 and 2 emissions divided by revenue. I divide emission intensity into four quantiles. In this table, I use scores to predict emissions one ($t + 1$) or two years ($t + 2$) ahead. Size is measured by capital or market capitalization as in table 2.1.

where $\exp(\varepsilon_{i,t})$ follows a gamma distribution with mean one for identification. I include *market cap* as a control, to account for the possibility that larger firms may receive more attention. *reporting score* is a pillar score (that contributes to the Environmental score just as the emission intensity scores) that measures the quantity and quality of environmental reporting. I include *reporting score* as a control, to account for the possibility that firms that report more have more to be examined. Since the correlation between emission intensity and scores is weak, as shown in previous facts, I include emission intensity not as a control, but as an informal validation check. For example, if emission intensity predicts misleading communication events but not emissions events, it suggests that RepRisk may not be reliable.

Regression results in table 2.4 show that firms with high scores will lead to more RepRisk events in the following year, after controlling for market capitalization and reporting scores. Furthermore, the coefficients on *reporting score* and *emission intensity* suggest that RepRisk reasonably captures events in these three categories. *reporting score* does not predict emissions events, but predicts misleading communication, whereas *emission intensity* does not predict misleading communication events, but predicts emissions events. In figure 2.2, I show the relationship between scores and RepRisk events. One can see that the distribution of RepRisk events is very skewed. Most companies have zero events, whereas those few companies that have high scores tend to have many events.

Fact 5: firms that are large, receive high scores, and have reported emissions at least once are associated with higher variance of documented emissions. I develop two ways to measure the variance of documented emissions. First, for a given firm i , I use the variance of the documented emissions of the firm in the entire sample over time. Second, I estimate the following time series model

$$\text{emission intensity}_{i,t+1} = \beta_{0,i} + \beta_1 \text{emission intensity}_{i,t} + \varepsilon_{i,t+1} \quad (2.5)$$

(2.5) is similar to an AR(1) model, but has the same AR(1) coefficient β_1 for all firms and the intercept $\beta_{0,i}$ serves as the firm fixed effect. I use the same β_1 for all firms due to the limited

	# RepRisk incidents next year		
	Misleading Comm.	Emissions	Environmental overall
	(1)	(2)	(3)
Env. score	0.011*** (0.003)	0.015*** (0.003)	0.013*** (0.003)
Env. reporting score	0.005** (0.002)	0.001 (0.002)	0.004** (0.002)
Market cap (log)	0.830*** (0.034)	0.856*** (0.036)	0.796*** (0.030)
Quintile 2	-0.075 (0.106)	0.091 (0.120)	0.245*** (0.093)
Quintile 3	-0.079 (0.100)	0.287** (0.113)	0.363*** (0.088)
Quintile 4	0.076 (0.103)	0.497*** (0.113)	0.497*** (0.094)
Quintile 5	-0.015 (0.100)	0.668*** (0.113)	0.829*** (0.099)
Time and Industry FE	Yes	Yes	Yes
BIC	7209.49	6545.83	11921.73
Observations	6,527	6,527	6,527

Note:

*p<0.1; **p<0.05; ***p<0.01
Market cap in \$ mil

Table 2.4: Regression results of (2.4).

time span of reported emissions. The variance of $\varepsilon_{i,t+1}$ is then the second way to measure reported emission variance.

$$\text{emissions variance}_i = \beta_0 + \beta_1 \text{predictor}_i + \varepsilon_i \quad (2.6)$$

I then use three predictors to predict emission variance: average firm size, average environmental score, and whether the firm has reported emissions at least once, over the sample period. Notice that since both measures of emission variance is computed over time, (2.6) is a cross-sectional regression (without the t subscript).


In table 2.5, I show that all three predictors are strong and statistically significant. I do not use all three predictors in the same regression because I want to avoid multicollinearity. As fact 1 suggests, high scores correlate strongly with size and whether having reported emissions. I compute the design matrix for the regression where all three predictors are included and the condition number is 31, greatly exceeding the threshold of 20 suggested by Belsley et al. (2005). The positive relationship is visualized in figure 2.3.

2.4. Model

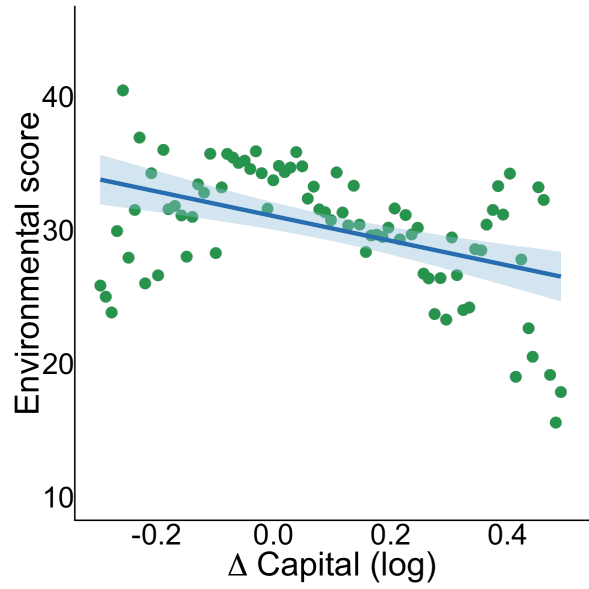
The key decision makers in my model are the atomistic firms. I will use a representative household to complete the general equilibrium model which, at the same time, helps me analyze welfare.

2.4.1. Firms

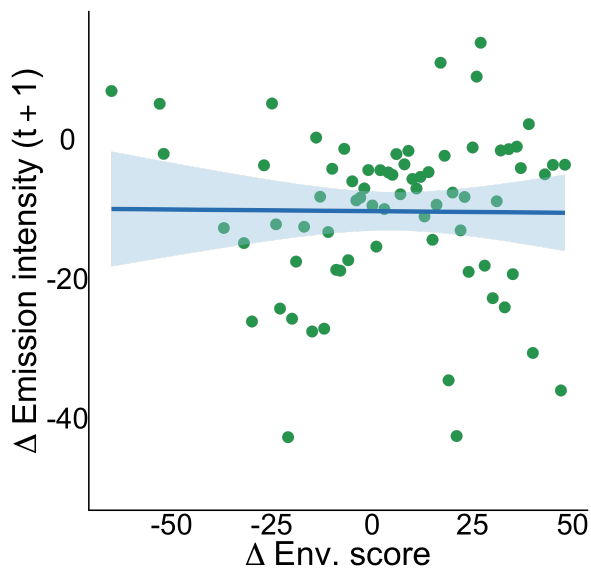
Time is continuous. There is a continuum of firms that differ in their green k_G and brown k_B capital. I define the greenness g of a firm as the ratio of green to aggregate capital: $g = k_G/k$, where $k = k_G + k_B$. Each firm's productivity A^n faces an idiosyncratic jump (independent of other firms) between two states $A^h > A^l$. Neither state is absorbing. This Poisson process is chosen for simplicity as in Achdou et al. (2022) and can be interpreted as the discrete analogue of a full-blown stochastic process as used in Barnett (2023). The intensity of this productivity jump is given by $\lambda^n \in \{\lambda^l, \lambda^h\}$.



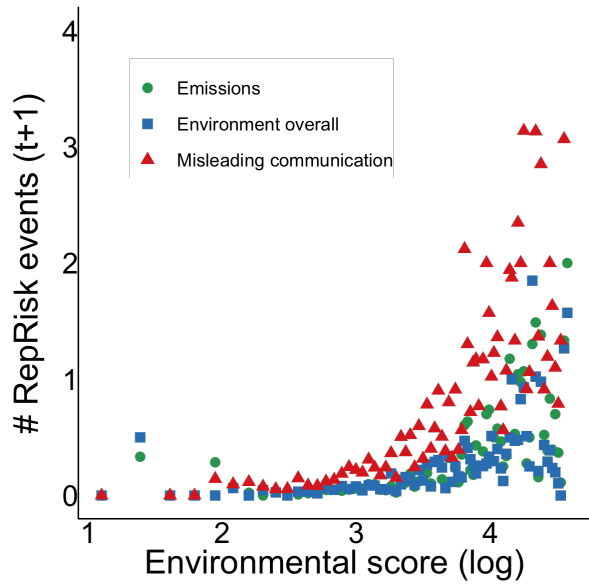
(a) Environmental score vs. capital



(b) Environmental score vs. change in capital



(c) Environmental score vs. emissions ($t + 1$)



(d) Environmental score vs. RepRisk events

Figure 2.2: Relationships between environmental scores, firm size, and emissions

	Emissions variance					
	Emissions var. (log)			Emissions residual var. (log)		
	(1)	(2)	(3)	(4)	(5)	(6)
Avg capital (log)	0.177** (0.072)			0.350*** (0.066)		
Avg env. score (log)		0.606*** (0.199)			0.745*** (0.217)	
Reported			0.541** (0.237)			1.086*** (0.164)
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,315	1,315	1,315	1,315	1,315	1,315
R ²	0.017	0.022	0.015	0.046	0.024	0.044

Note: *p<0.1; **p<0.05; ***p<0.01
Capital in \$ mil

Table 2.5: Regression results of (2.6). Here, *Reported* is a dummy variable for whether the firm has reported emissions at least once over the sample period.

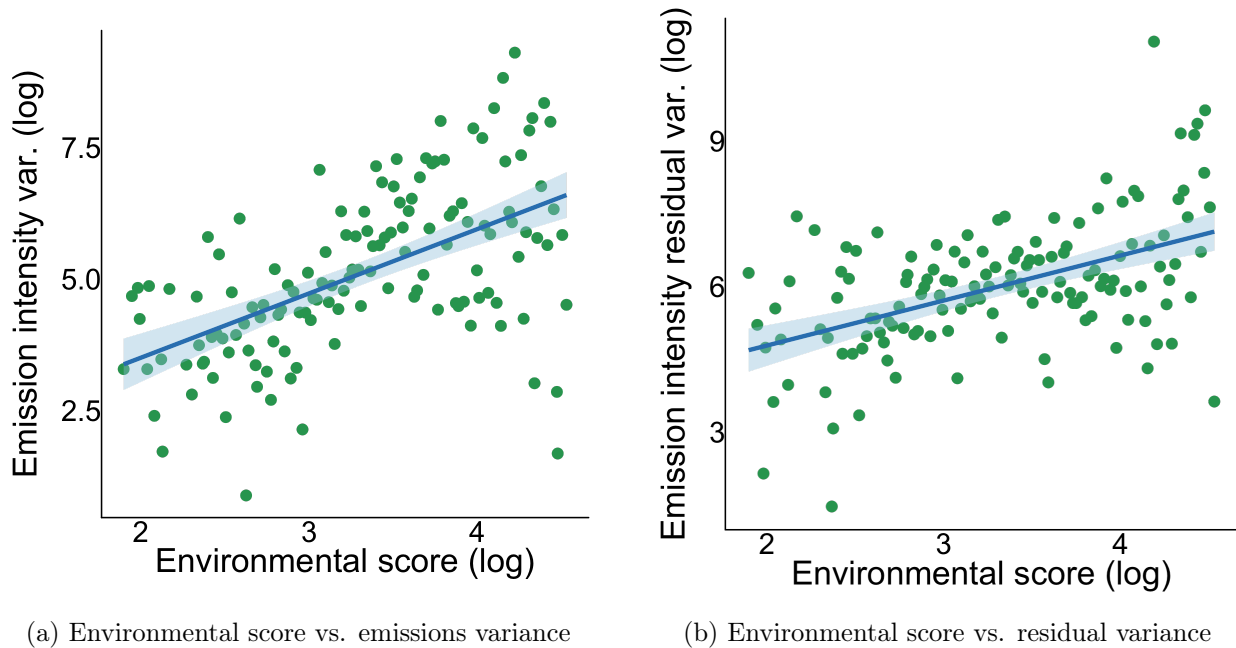


Figure 2.3: Relationships between environmental scores and emission variability

$$y(k_{Gt}, k_{Bt}, A_t^n, \mathcal{S}_t) = A_t^n D(\mathcal{S}_t) \left((k_{Gt}^{\alpha_G})^{\frac{\alpha-1}{\alpha}} + (k_{Bt}^{\alpha_B})^{\frac{\alpha-1}{\alpha}} \right)^{\frac{\alpha}{\alpha-1}} \quad (2.7)$$

$D(\mathcal{S}_t)$ represents climate risk. I assume that $D(\mathcal{G}) = 1$ and $D(\mathcal{B}) \in (0, 1)$. \mathcal{S}_t can take only two values: \mathcal{G} and \mathcal{B} . \mathcal{S}_t denotes the climate state that the economy is in. The economy starts out in the good (\mathcal{G}) state. With some intensity $\lambda^{\mathcal{S}}$, it jumps to the bad (\mathcal{B}) state permanently so that the bad state is absorbing. In other words, I model climate risk as a Poisson jump process that scales down the production of *all firms*. Climate risk is aggregate and will therefore be priced. This parsimonious way of modeling climate risk is informed by Hong et al. (2023a) and Lenton et al. (2008). $D(\mathcal{B})$ is a constant and thus left as a parameter to calibrate (more in section 2.5). Note that unlike climate risk, the idiosyncratic jump in productivity (between A^h and A^l) is not priced. The functional form in (2.7) is informed by Acemoglu et al. (2012) and Acemoglu et al. (2017), the latter of which shows firms have incentives to diversify their energy portfolio across different energy sources.

Brown capital produces emissions. Firms must report their emissions. The evolution of *reported* emissions, E_t^R , in the absence of greenwashing, is given by

$$dE_t^R = \mu_{k_B} \eta k_t (1 - g_t) dt + \sigma_B \eta (1 - g_t) k_t dW_t \quad (2.8)$$

(2.8) is motivated by a commonly used functional form, which models emissions linearly as $\eta k_t (1 - g_t)$ and is based on the IPCC report (Bruckner et al., 2014) which finds that emissions are proportional to fossil fuels (with the proportionality constant being η). The approximation is widely used in the literature, such as in Barnett et al. (2023a), Barnett et al. (2020), and Karydas and Xepapadeas (2019). My addition of dW_t , i.i.d. across firms, is to allow for measurement error in emissions and to introduce a source of uncertainty. Formally, $\{W_t\}_{t \geq 0}$ is a Brownian motion under a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. σ_B and η are left as parameters to calibrate (more in section 2.5). The productivity jumps, climate jump, and randomness in emissions are the only source of risk in my model, but the climate jump is the only priced risk.

Going forward, I suppress t for notational convenience whenever possible. Firms invest in brown and green capital subject to quadratic adjustment costs. Denote a firm's investment in capital j as ι_j . Then the evolution of each capital type j is given by

$$dk_j/k_j = \underbrace{(\iota_j - \delta)}_{\mu_{k_j}} dt \quad (2.9)$$

where δ is the depreciation rate. By Ito's lemma, the greenness g and aggregate capital k evolve according to

$$\begin{aligned} dg/g &= \underbrace{(1-g)(\iota_G - \iota_B)}_{\mu_g} dt \\ dk/k &= (\iota_G g + \iota_B(1-g) - \delta) dt \end{aligned} \quad (2.10)$$

The idiosyncratic jump in productivity, together with adjustment costs (introduced later), gives rise to a distribution of firms. The key friction in the model is that greenness g is not observed and the mapping between emissions and greenness is clouded by randomness (dW_t). The only observable state variables to households are A^n (productivity), \mathcal{S} (climate state), and k (aggregate capital). Note that since there is a distribution of firms, households cannot infer greenness based on the aggregates.

Carbon tax. In the bad \mathcal{B} state, in addition to the loss in production, all firms face a carbon tax proportional to their brown capital: firms pay $\tau\eta k_{Bt}$ in the bad state with τ being the tax rate. This tax on brown capital is equivalent to an upstream carbon tax (Weisbach and Metcalf, 2009). Emissions cannot be directly observed, but the government can tax brown capital (e.g. fossil fuels). This approach effectively targets carbon emissions at their source by taxing the input materials that generate them. Greener firms (higher g_t) are less affected by the carbon tax. Together with the loss in production that affects all firms, the carbon tax will give rise to a greenium: greener firms have a higher cash flow when production (and hence consumption) is low. Greener firms, in the good state, have a lower cost of capital because they are hedges against climate risk. Households therefore have

an incentive to figure out g (described below). A negative greenium is motivated by the empirical literature. For example, Choi et al. (2020) provide evidence that when investors are concerned about climate change, carbon-intensive stocks underperform. Pastor et al. (2022) estimate a lower expected return and implied cost of capital for firms with high ESG scores. Furthermore, I assume that the carbon tax is visible to the households and subsequently g is correctly inferred in the bad \mathcal{B} state.

Firms' HJB equation in the bad \mathcal{B} state. When $\mathcal{S}_t = \mathcal{B}$, transition risk is realized. There is no aggregate risk. Therefore, the required rate of return for firms is the risk-free rate r_t . Furthermore, I assume that households can see the carbon tax and correctly infer g_t in the bad \mathcal{B} state. The cash flow in the bad state is given by

$$\begin{aligned} \text{CF}(g_t, k_t, A_t^n, \mathcal{B}_t) = & A_t^n D(\mathcal{B}_t) \left((k_{Gt}^{\alpha_G})^{\frac{\alpha-1}{\alpha}} + (k_{Bt}^{\alpha_B})^{\frac{\alpha-1}{\alpha}} \right)^{\frac{\alpha}{\alpha-1}} \\ & - \tau \eta k_{Bt} - \sum_{j \in \{G, B\}} \left(\iota_{jt} + \frac{\theta_j}{2} \iota_{jt}^2 \right) k_{jt} \end{aligned} \quad (2.11)$$

Firm value $Q(g, k, A^n, \mathcal{B})$ in the bad \mathcal{B} state satisfies the following HJB equation:

$$rQ(g, k, A^n, \mathcal{B}) = \max_{\iota_B, \iota_G} \text{CF}(g, k, A^n, \mathcal{B}) + \mu_k k Q_k + \mu_g g Q_g + \lambda^n (Q(A^{-n}, \mathcal{B}) - Q(A^n, \mathcal{B})) \quad (2.12)$$

Note that the firm value in the bad \mathcal{B} state is a function of g , k , and A^n , not \hat{g} .

2.4.2. Households

Households have log utility:

$$u(C_t, N_t) = \rho \ln(C_t/N_t)$$

where N_t is climate damages, elaborated in section 2.4.3. Her consumption is then given by her discount rate ρ multiplied by her wealth. Let \mathbb{M}_t denote the stochastic discount factor (SDF) of the household. The SDF then follows

$$\frac{d\mathbb{M}_t}{\mathbb{M}_t} = -r_t dt - j_t^C (dJ_t - \lambda^S dt) \quad (2.13)$$

where r_t is the risk-free rate (which is equal to ρ under log utility) and dJ_t denotes the climate Poisson jump with intensity λ^S . j_t^C is the drop in consumption after the climate jump, determined in equilibrium (given by (2.29)).

Learning. Households can observe E_t but not g_t . In the bad \mathcal{B} state, households can correctly compute g_t through the carbon tax. In the good \mathcal{G} state, households' estimation of g_t is given as follows. I assume that households use the simplified (and misspecified) model below, a Kalman-Bucy filter, to approximate the stochastic process $\{E_t^R : t \geq 0\}$:

$$\begin{aligned} d\tilde{k}_{Bt} &= \tilde{\mu}_{k_B} \tilde{k}_{Bt} dt \\ dE_t^R &= \eta \tilde{k}_{Bt} dt + \eta \sigma_E dW_t \end{aligned} \quad (2.14)$$

where E_t^R is the reported emissions, the true stochastic process of which is given by (2.8), and $\tilde{\mu}_{k_B}$ is a parameter (more in section 2.5). Concretely, \tilde{k}_{Bt} is the latent state in the Kalman-Bucy filter, the posterior distribution of which households aim to estimate. E_t^R is the observed signal in the Kalman-Bucy filter. σ_E is a constant (i.e. a parameter). Let \hat{k}_{Bt} denote the posterior mean of \tilde{k}_{Bt} . One can show that the posterior mean of the Kalman-Bucy filter follows (Øksendal, 1998)

$$\begin{aligned} d\hat{k}_{Bt} &= \left(\tilde{\mu}_{k_B} - \frac{\eta^2 S_t}{\sigma_E^2} \right) \hat{k}_{Bt} dt + \frac{\eta S_t}{(\eta \sigma_E)^2} dE_t^R \\ \frac{dS_t}{dt} &= 2\tilde{\mu}_{k_B} S_t - \frac{\eta^2}{(\eta \sigma_E)^2} S_t^2 \end{aligned}$$

As a further simplification, notice for sufficiently large t , S_t converges to its mean, given by

$$S^* = 2\tilde{\mu}_{k_B} \sigma_E^2$$

and I can rewrite, by substituting S^* into the Kalman-Bucy filter,

$$d\widehat{k}_{Bt} = -\widetilde{\mu}_{k_B}\widehat{k}_{Bt}dt + \frac{2\widetilde{\mu}_{k_B}}{\eta}dE_t^R$$

That is, I assume that households use the long-run variance of E_t to estimate k_B . Furthermore, define $\widehat{g}_t = 1 - \frac{\widehat{k}_{Bt}}{k_t}$. \widehat{g}_t is the estimated greenness of the firm. I assume that households use the posterior mean of \widetilde{k}_B as their estimate of k_B , and at any given time they set $k_B = \widehat{k}_B$ and $g = \widehat{g}$.

Next, by substituting (2.8) into the Kalman-Bucy filter (and Ito's lemma), I have

$$d\widehat{k}_{Bt} = \left(-\widetilde{\mu}_{k_B}\widehat{k}_{Bt} + 2\mu_{k_B}\widetilde{\mu}_{k_B}k_{Bt}\right) dt + 2\mu_{k_B}\sigma_B k_{Bt}dW_t$$

which can be rewritten as

$$\frac{d\widehat{k}_{Bt}}{\widehat{k}_{Bt}} = \left(-\widetilde{\mu}_{k_B} + 2\mu_{k_B}\widetilde{\mu}_{k_B}\frac{1-g}{1-\widehat{g}}\right) dt + 2\mu_{k_B}\sigma_B\frac{1-g}{1-\widehat{g}}dW_t \quad (2.15)$$

with

$$\mu_{\widehat{k}_B} = -\widetilde{\mu}_{k_B} + 2\mu_{k_B}\widetilde{\mu}_{k_B}\frac{1-g}{1-\widehat{g}} \quad \text{and} \quad \sigma_{\widehat{k}_B} = 2\mu_{k_B}\sigma_B\frac{1-g}{1-\widehat{g}}$$

Applying Ito's lemma, I obtain

$$d\left(\frac{\widehat{k}_{Bt}}{k_t}\right) = (1-\widehat{g})\left[\left(\mu_{\widehat{k}_B} - \mu_k\right) dt + \sigma_{\widehat{k}_B} dW_t\right]$$

Then the posterior mean of greenness, \widehat{g} , is given by (appendix A.2)

$$d\widehat{g} = (\widehat{g} - 1)\left[\left(\mu_{\widehat{k}_B} - \mu_k\right) dt + \sigma_{\widehat{k}_B} dW_t\right]$$

In summary, when households adopt model (2.14), the posterior mean of the Kalman-Bucy filter is

given by

$$d\hat{g}_t = (\hat{g} - 1) \left[\left(\mu_{\hat{k}_B} - \mu_k \right) dt + \sigma_{\hat{k}_B} dW_t \right] \quad (2.16)$$

Distorted learning and greenwashing. Firms may adopt different models for the brownian noise dW_t . Therefore, instead of adopting model (2.8), they may choose a different probability measure. Formally, instead of probability measure \mathbb{P} , they choose a probability measure $\tilde{\mathbb{P}}$. I restrict my analysis to the set of probability measures $\tilde{\mathbb{P}}$ under which $d\tilde{W}_t$ is a Brownian motion. Courtesy of the Girsanov theorem (Hansen and Sargent (2011), Hansen et al. (2019), Hansen and Sargent (2022)), this is equivalent to a drift distortion:

$$d\tilde{W}_t = x_{Et} dt + dW_t$$

If firms use $d\tilde{W}_t$ instead of dW_t , the posterior mean of the Kalman-Bucy filter is given by

$$d\hat{g}_t = (\hat{g} - 1) \left[\left(\mu_{\hat{k}_B} - \mu_k + \sigma_{\hat{k}_B} x_{Et} \right) dt + \sigma_{\hat{k}_B} dW_t \right] \quad (2.17)$$

(2.17) is obtained by taking (2.16) and replacing dW_t with $x_{Et} + dW_t$.

I assume that firms greenwash by choosing x_{Et} at a quadratic cost, described in (2.25). Therefore, I model greenwashing as a distortion in the reporting (and subsequently learning) process that firms pay for.

One may interpret x_{Et} (or equivalently a different probability measure) as different reporting standards, carbon accounting methods, and emission model. This is motivated by facts 1 and 5 in section 2.3: firms that report emissions must choose an emission model, as emissions cannot be directly observed or recorded. These firms receive a higher score. Distorting the probability measure implies a departure from the normal distribution with mean zero and variance σ_B^2 , possibly leading to higher variance in the emission noise. Furthermore, this is grounded in reality. In the CSA sent by S&P Global, firms may choose different reporting standards. The SEC does not have any

regulations on emissions reporting. In short, changing x_{Et} can be interpreted as choosing a different measurement technique.

I assume that households mistakenly (i.e. as a friction) behave as if there is no greenwashing (i.e. they behave as if $x_{Et} = 0$ for all firms). This assumption does not depart too much from reality. Greenwashing has only become a concern in recent years. Even so, there is scant clarity on greenwashing penalties or emissions reporting standards. The responses to S&P Global's CSA are confidential. Investors are constrained in their ability to detect greenwashing.

Required rate of return. Since households assume $g = \hat{g}$ and no greenwashing, they behave as if firms' cash flow is given by

$$\begin{aligned} \text{CF}_t^H(\hat{g}, k, A^n, \mathcal{S}) &= A_t^n D(\mathcal{S}_t) \left((k_{Gt}^{\alpha_G})^{\frac{\alpha-1}{\alpha}} + (\hat{k}_{Bt}^{\alpha_B})^{\frac{\alpha-1}{\alpha}} \right)^{\frac{\alpha}{\alpha-1}} \\ &\quad - 1(\mathcal{S}_t = \mathcal{B}) \tau \hat{k}_{Bt} - \left(\iota_{Gt} + \frac{\theta_G}{2} \iota_{Gt}^2 \right) k_{Gt} - \left(\iota_{Bt} + \frac{\theta_B}{2} \iota_{Bt}^2 \right) \hat{k}_{Bt} \end{aligned} \quad (2.18)$$

In the good state, and under the household belief, $\mathbb{M}_t \text{CF}_t^H dt + d(\mathbb{M}_t Q_t^H)$, where Q_t^H is the firm value under the households' belief, follows a martingale (Duffie, 2010). Setting the drift of $\mathbb{M}_t \text{CF}_t^H dt + d(\mathbb{M}_t Q_t^H)$ to zero, I obtain the following equality:

$$r + \lambda^S j^S j^C = \max_{\iota_B, \iota_G} \frac{1}{Q^H} [\text{CF}^H(\mathcal{G}) + Q_k^H \mu_k + Q_g^H \mu_g(\hat{g})\hat{g}] + \lambda^n \left(\frac{Q^H(A^{-n}, \mathcal{G})}{Q^H(A^n, \mathcal{G})} - 1 \right) \quad (2.19)$$

where j^S , the price of jump risk, is given by

$$j^S = \frac{Q(\hat{g}, k, A^n, \mathcal{B})}{Q^H(\hat{g}, k, A^n, \mathcal{G})} - 1 \quad (2.20)$$

$Q(\hat{g}, k, A^n, \mathcal{B})$ is the value function of the firm as described in (2.12), evaluated at \hat{g} . In the bad state, households no longer form beliefs about the firm's greenness and there is no greenwashing or greenium, so $Q^H(\mathcal{B}, \hat{g}) = Q(\mathcal{B}, g)$ for all $\hat{g} = g$.

Using the martingale approach in Brunnermeier and Sannikov (2016), I can obtain the required rate

of return for firms:

$$\hat{r}(\hat{g}, k, A^n) = r + \lambda^S j^S j^C \quad (2.21)$$

Note that the required rate of return \hat{r} depends on \hat{g} but not g , because j^S is a function of \hat{g} . (2.20) shows the source of greenium and the incentive to greenewash. Since production drops after the climate jump and consumption is proportional to wealth, I have $j^C < 0$. Firms' value drops after the climate jump as well, so $j^S < 0$. Together, I have $j^S j^C > 0$. The less negative j^S is, the lower cost of capital the firm has. Greener firms pay a smaller carbon tax after the climate jump so their post-jump value function is higher. This gives rise to the greenium: households want to hold firms with a higher g , and an incentive to greenwash: firms want a higher \hat{g} .

2.4.3. Climate

I use the following process for temperature anomaly T_t :

$$dT_t = \beta_f (E_t + \bar{E}) dt \quad (2.22)$$

where β_f is the climate sensitivity parameter that translates emissions to temperature. This specification is motivated by Matthews et al. (2009), commonly used in the literature (Barnett et al. (2023a), Brock and Xepapadeas (2017) and Matthews et al. (2009)). \bar{E} is the emissions that are not generated by the firms. In my model, \bar{E} is exogenous and will be calibrated to the average emissions from 2010 to 2020 outside of the United States, since my model will be calibrated to the United States. E_t is the aggregate emissions generated by all the firms. Let K_{Bt} denote aggregate brown capital. I assume that

$$E_t = \eta K_{Bt} \quad (2.23)$$

Climate damages evolve according to

$$d \log N_t = (\gamma_1 + \gamma_2 T_t) \beta_f (E_t + \bar{E}) dt \quad (2.24)$$

Notice that climate damages N_t are an externality. Households and firms have no control over N_t .

2.4.4. Equilibrium

Firms' HJB equation in the good \mathcal{G} state. Since households set $g = \hat{g}$, the cost of capital depends on \hat{g} instead of g . Firms have an incentive to distort the learning of households by choosing x_{Et} to have a lower cost of capital. Nevertheless, distortion is not free. Firms pay a cost $\frac{\theta_E}{2} x_{Et}^2$ for greenwashing. One can interpret this cost as fees spent on ESG consulting and potential repurchases from legal woes and bad press. This is motivated by fact 4: firms with high scores do receive more bad press.

The cash flow of a firm in the good \mathcal{G} state is given by

$$\begin{aligned} \text{CF}_t(g, k, A^n, \mathcal{G}) &= A_t^n D(\mathcal{G}_t) \left((k_{Gt}^{\alpha_G})^{\frac{\alpha-1}{\alpha}} + (k_{Bt}^{\alpha_B})^{\frac{\alpha-1}{\alpha}} \right)^{\frac{\alpha}{\alpha-1}} \\ &\quad - \sum_{j \in \{G, B\}} \left(\iota_{jt} + \frac{\theta_j}{2} \iota_{jt}^2 \right) k_{jt} - \frac{\theta_E}{2} x_{Et}^2 \end{aligned} \quad (2.25)$$

Firms take the required rate of return \hat{r} as given. The state variables of the firm are k (aggregate capital), g (greenness), \hat{g} (perceived greenness), \mathcal{S} (climate state), and A^n (productivity). Firm value $Q(g, k, \hat{g}, A^n, \mathcal{G})$ in the good \mathcal{G} state must satisfy the following HJB equation:

$$\begin{aligned} \hat{r}(\hat{g}, k, A^n) Q(g, k, \hat{g}, A^n, \mathcal{G}) &= \max_{\iota_B, \iota_G, x_E} \text{CF}(\mathcal{G}) \\ &\quad + \mu_k k Q_k + \mu_g g Q_g + (\hat{g} - 1) \left(\mu_{\hat{k}_B} - \mu_k + \sigma_{\hat{k}_B} x_{Et} \right) Q_{\hat{g}} + \frac{1}{2} \left(\frac{\sigma_{\hat{k}_B}}{k} \right)^2 Q_{\hat{g}\hat{g}} \\ &\quad + \lambda^{\mathcal{S}} (Q(A^n, \mathcal{B}) - Q(A^n, \mathcal{G})) + \lambda^n (Q(A^{-n}, \mathcal{G}) - Q(A^n, \mathcal{G})) \end{aligned} \quad (2.26)$$

Stationary distribution and market clearing. To determine j_t^C , the drop in consumption after the climate jump *under households' belief*, I need to solve for the stationary distribution of the firms under households' belief, since the representative household's consumption is equal to the aggregate consumption, which in turn depends on the stationary distribution of the firms. Since the households cannot conceive of greenwashing and believe that $\hat{g} = g$, the Kolmogorov forward

equation is given by

$$0 = -\frac{d}{dk} [M_{\ln k} \mu_{\ln k}] - \frac{d}{d\hat{g}} [M_g \mu_g(\hat{g}) \hat{g}] - \lambda^n M(A^n) + \lambda^n M(A^{-n}) \quad (2.27)$$

for $\mathcal{S} \in \{\mathcal{G}, \mathcal{B}\}$. Let $M(k, \hat{g}, A^n, \mathcal{S})$, the solution to (2.27), be the stationary distribution.

The market clearing conditions are given as follows. Let ι_B^H, ι_G^H be the optimal controls to (2.19), i.e. firms' optimal investment rates under households' belief. For notational convenience, define the following terms:

$$\begin{aligned} I_B(\mathcal{S}) &= \int \iota_B^H (1 - \hat{g}) k dM(k, \hat{g}, A^n | \mathcal{S}) && \text{(Aggregate brown investment)} \\ I_G(\mathcal{S}) &= \int \iota_G^H \hat{g} k dM(k, \hat{g}, A^n | \mathcal{S}) && \text{(Aggregate green investment)} \\ I(\mathcal{S}) &= I_B + I_G && \text{(Total investment)} \\ Y(\mathcal{S}) &= \int A^n D(\mathcal{S}) \left((k_G^{\alpha_G})^{\frac{\alpha-1}{\alpha}} + (k_B^{\alpha_B})^{\frac{\alpha-1}{\alpha}} \right)^{\frac{\alpha}{\alpha-1}} \\ &\quad \times dM(k, \hat{g}, A^n | \mathcal{S}) && \text{(Aggregate output)} \\ \mathcal{T}(\mathcal{S}) &= \int 1(\mathcal{S} = \mathcal{B}) \tau \eta (1 - \hat{g}) k dM(k, \hat{g}, A^n | \mathcal{S}) && \text{(Aggregate tax revenue)} \end{aligned}$$

Market clearing requires that, given each \mathcal{S} ,

$$C = Y - I + \mathcal{T}$$

where C is aggregate consumption. Log utility implies that

$$C = \rho \mathcal{W}$$

so I must have

$$\rho \mathcal{W} = Y - I + \mathcal{T} \quad (2.28)$$

The consumption jump j^C is given by

$$j^C = \frac{\mathcal{W}(\mathcal{B})}{\mathcal{W}(\mathcal{G})} - 1 \quad (2.29)$$

Notice for each firm, the required rate of return \hat{r} depends on j_t^C , which in turn depends on the stationary distribution of the firms. The firms, in other words, are atomistic players in a *mean-field game* (Lasry and Lions, 2007).

Equilibrium concept. I will characterize a mean-field game (MFG) equilibrium. In this MFG equilibrium, the required rate of return $\hat{r}(k, \hat{g}, A^n)$ is determined by the stationary distribution $M(k, \hat{g}, A^n, \mathcal{S})$ of the firms, *under households' belief*. Given $\hat{r}(k, \hat{g}, A^n)$, the policy functions $\{\iota_G, \iota_B, x_E\}$ are determined and the value function $Q(k, g, \hat{g}, A^n, \mathcal{S})$ satisfies the HJB equations (2.26) and (2.12).

A mean field game equilibrium is given by value function $Q(k, g, \hat{g}, A^n, \mathcal{S})$, stationary distribution $M(k, \hat{g}, A^n, \mathcal{S})$, required rates of return $\hat{r}(k, \hat{g}, A^n)$, and policy functions $\iota_G(k, g, \hat{g}, A^n, \mathcal{S})$, $\iota_B(k, g, \hat{g}, A^n, \mathcal{S})$, $x_E(k, g, \hat{g}, A^n, \mathcal{S})$ such that

- Given the policy functions, stationary distribution, and required rate of return, value function Q solves the HJB equations associated with the firms' problem (2.26) and (2.12);
- Given the value function, stationary distribution, and required rates of return, each firm cannot obtain a higher payoff through unilateral deviations from the policy functions ι_G, ι_B, x_E ;
- Given the required rates of return, stationary distribution M solves the Kolmogorov Forward Equation (2.27);
- Given the value function, stationary distribution, and policy functions, the required rate of return \hat{r} satisfies (2.21) and market clearing. In particular, the jump in consumption j^C is given by (2.29).
- The policy functions ι_G, ι_B, x_E solve the firms' problem (2.26) and (2.12).

Mean-field games are not new in economics. For example, Achdou et al. (2022) use a mean-field game to study inequality questions. Alvarez et al. (2022) use a mean-field game to model price setting. My definition of a mean-field game equilibrium follows the standard definition of mean-field games. Furthermore, the equilibrium concept in definition 2.4.4 is not a perfect Bayesian equilibrium, because households cannot conceive of greenwashing and believe that $\hat{g} = g$.

Solving for this equilibrium is computationally intensive. In practice, I solve for the value functions in logs, the derivation of which is provided in the appendix A.1. The numerical algorithm is detailed in section A.3.

2.4.5. Social cost of carbon

I characterize the social planner problem so that I can quantify the welfare consequences of greenwashing through the social cost of carbon. Let K_t be the aggregate capital in the economy and C_t, K_{Bt}, K_{Gt}, G_t similarly defined. I assume that the planner is subject to the same productivity and climate jumps. The state variables of the planner are $K, G, A^n, \mathcal{S}, T, N$. The planner's problem is given by

$$\begin{aligned}
& \max_{\{C_t, \iota_{Gt}, \iota_{Bt}\}_{t=0}^{\infty}} \mathbb{E} \left[\int_{t=0}^{\infty} \exp(-\rho t) U(C_t, N_t) dt \right] \\
& \text{subject to } C_t + \left(\iota_{Gt} + \frac{\theta_G}{2} \iota_{Gt}^2 \right) K_{Gt} + \left(\iota_{Bt} + \frac{\theta_B}{2} \iota_{Bt}^2 \right) K_{Bt} = Y_t(\mathcal{S}_t) \\
& \quad Y_t(\mathcal{S}_t) = A_t^n D(\mathcal{S}_t) \left((K_{Gt}^{\alpha_G})^{\frac{\alpha-1}{\alpha}} + (K_{Bt}^{\alpha_B})^{\frac{\alpha-1}{\alpha}} \right)^{\frac{\alpha}{\alpha-1}} \\
& \quad dK_{Gt}/K_{Gt} = (\iota_{Gt} - \delta) dt \\
& \quad dK_{Bt}/K_{Bt} = (\iota_{Bt} - \delta) dt \\
& \quad dT_t = \beta_f (E_t + \bar{E}) dt \\
& \quad d \ln N_t = (\gamma_1 + \gamma_2 T_t) \beta_f (E_t + \bar{E}) dt
\end{aligned} \tag{2.30}$$

Let $V(K, G, A^n, \mathcal{S}, T, N)$ be the value function of the planner. I postulate that the value function takes the following form: $V(K, G, A^n, \mathcal{S}, T, N) = v(K, G, T, A^n, \mathcal{S}) - \ln N$. $v(K, G, T, A^n, \mathcal{S})$ must

satisfy the following HJB equation:

$$\begin{aligned}
\rho v(K, G, T, \mathcal{S}) = & \max_{\{C, \iota_G, \iota_B\}} \rho \ln C + v_{\ln K} (\iota_G G_t + \iota_B (1 - G_t) - \delta) + v_G G(1 - G) (\iota_G - \iota_B) \\
& + v_T (\beta_f (E_t + \bar{E})) - ((\gamma_1 + \gamma_2 T_t) \beta_f (E_t + \bar{E})) + \lambda^{\mathcal{S}^-} (v(\mathcal{S}^-) - v(\mathcal{S})) \\
& + \lambda^n (v(A^n, \mathcal{S}) - v(A^n, \mathcal{S}))
\end{aligned} \tag{2.31}$$

I follow the approach that Barnett et al. (2020) and Cai and Lontzek (2019) use to compute the social cost of carbon: marginal utility of emissions scaled by the marginal utility of consumption. The HJB equation (2.31) implies that the social cost of carbon is given by

$$\text{scc}_t = \frac{C^*}{\rho} \left[v_T \beta_f - (\gamma_1 + \gamma_2 T) \beta_f \right] \tag{2.32}$$

with C^* being the optimal consumption, i.e. the optimal control of (2.30) and (2.31).

2.5. Calibration

I summarize the calibrated parameters in table 2.6. Most of the parameters are taken from the literature. For example, the productivity parameters (both levels and jump intensities, $A_h, A_l, \lambda_h, \lambda_l$) are based on the RBC literature. To approximate the $AR(1)$ process for productivity with a two-state Markov chain, I use the method in Rouwenhorst (1995), which requires that I provide the long run mean of productivity. To find the mean of productivity, I choose the long run mean of productivity such that the mean capital is \$17 billion, i.e. $\mathbb{E}[k] = 17$. Climate disaster parameters ($D(B), \lambda^{\mathcal{S}}$) are taken from Hong et al. (2023a). Parameters such as discount rate ρ and depreciation rate δ are taken from Barnett et al. (2023a) and are in line with the literature. The carbon tax parameter τ is taken from Fried et al. (2021) and is in the range informed by the Cap-and-Trade program in California, which provides traded prices of carbon allowances. Moreover, it is supported by the recommendations in IMF (2019). I use the emission parameter η and damage parameters γ_1, γ_2 from Barnett et al. (2023a). I calibrate $\bar{E} = 29$ because the average emissions from the rest of the world is 29 gigatons from 2012 - 2022. I choose $\sigma_B = 0.08$ by computing the variance in documented emissions provided by Trucost. I choose $\tilde{\mu}_{k_B} = 0.03$ because emissions have grown at

3% on average in the past decade.

The remaining parameters are θ_G and θ_E . I choose θ_G such that the pre-climate jump economy has average greenness of $\mathbb{E}[g] = 0.67$, a choice close to Carattini et al. (2023). I calibrate these two parameters by matching the moments of a model where households can perfectly observe greenness. θ_E is left as a free parameter that I vary to see how it induces greenwashing.

2.5.1. Validation

I verify the empirical validity of my calibration by comparing the moments in an economy without greenwashing to the literature. In table 2.7, I summarize the results in the economy where households can perfectly observe greenness g . The model does a good job at matching the moments. Furthermore, quantitative implications are in line with the literature: first, a -4% drop in production and consumption after jumping to the bad state agrees with the estimates in Desmet and Rossi-Hansberg (2023), who claim that climate change may reduce global GDP by between 2% and 6%. Second, my model predicts that a carbon tax of \$45 reduces emissions by 51%. Carattini et al. (2023) estimate a drop of 40% in emissions from a \$17 carbon tax, whereas IMF (2019) estimates a drop of more than 20% from a \$50 carbon tax. While empirical estimates are inconclusive, based on the theoretical results in the literature, the model makes sensible predictions. Third, the average firm value drops by 0.7% after the climate jump. While literature and data offer no benchmark, Hong et al. (2023a) provide empirical evidence that in the presence of climate disasters, the real interest rate and equity premium drop by 0.09% and 0.31% respectively. Moreover, Pastor et al. (2022) estimate that “greennium” is about 1%. Despite being different measures, these estimates confirm that the expected firm value change implied by my model is not unexpected. The low increase (13%) in greenness is due to the high green adjustment cost (θ_G), highlighting the importance of R&D which should either reduce adoption cost or raise productivity (Barnett et al., 2023a). Overall, the baseline model without greenwashing generates implications that are consistent with the literature and data.

Table 2.6: Calibrated parameters

Description	Choice	Target/Source
Discount rate (ρ)	0.00375	Barnett et al. (2020)
Depreciation (δ)	0.05	Barnett et al. (2023a)
Green adjustment cost (θ_G)	95.8	Calibrated
Brown adjustment cost (θ_B)	10	Barnett et al. (2023a)
Productivity, high (A_h)	25.29	RBC literature, own calibration
Productivity, low (A_l)	7.02	RBC literature, own calibration
Productivity jump intensity (λ_h)	0.025	RBC literature
Productivity jump intensity (λ_l)	0.025	RBC literature
Green capital intensity (α_G)	0.35	Carattini et al. (2023)
Brown capital intensity (α_B)	0.35	Carattini et al. (2023)
CES parameter (α)	2	Carattini et al. (2023)
Greenwashing cost (θ_E)	450	Calibrated
Emissions variance (σ_B)	0.08	Data
Mean brown capital growth ($\tilde{\mu}_{k_B}$)	0.03	Data
Emission (η)	0.0192	Calibrated
Carbon tax (τ)	45	Fried et al. (2021), Cap-and-trade
Climate damages (D)	0.978	Hong et al. (2023a)
Climate jump intensity (λ_S)	0.02	Hong et al. (2023b)
Climate damages (γ_1, γ_2)	0.00017675, 0.3377	Barnett et al. (2023a)
Exogenous emissions (\bar{E})	29	Data

Note: parameters are calibrated on an annual basis.

	Description	Model	Data/literature
Moments	Investment ratio $\mathbb{E}[\iota_G + \iota_B]$	0.10	0.10
	Greenness ($\mathbb{E}[g]$)	0.67	0.67
	Average capital ($\mathbb{E}[k]$)	17	17
Implications	Consumption change (j^C)	-5.68	[-6, -2]
	Production change	-5.89	[-6, -2]
	Emissions change	-51.71	[-40, -20]
	Greenness change	0.13	Not available
	Firm value change (j^S)	-0.7	Not available

Note: expectations are taken over the stationary distribution of the firms M .

Table 2.7: Baseline results: first half contains untargeted and **targeted** moments; second half contains quantitative implications (change in percents when the economy jumps to the bad state) from the calibration in section 2.5.

2.6. Quantitative Analysis

Next, I introduce greenwashing, allowing x_{Et} to differ from zero. I introduce an “MIT” shock, where the economy transitions from one with no greenwashing to one with greenwashing. I simulate $N = 200$ firms with their capital fixed at the mean and greenness fixed at 0.3, and compute the transition paths of different variables of interest (greenness g , estimated greenness \hat{g} , capital $\ln k$, and emissions E), for different values of $\theta_E \in \{30, 50\}$, over 10 years. Since θ_E is the penalty on greenwashing, higher θ_E implies greenwashing is more expensive for the firms. In addition, I simulate the transition paths of models where households can perfectly observe firms’ greenness (“Perfect information”) and where households cannot observe firms’ greenness and firms cannot greenwash (“No greenwashing”).

The results are shown in figure 2.4, where I plot the average value of the firms in each month for different model specifications. First, when firms can perfectly observe greenness (“Perfect information”), the economy is greener. As the penalty on greenwashing (θ_E) decreases, greenness declines as well, highlighting the incentive to greenwash. Estimated greenness \hat{g} (figure 2.4b) shows a reversed pattern. When θ_E is low (i.e. greenwashing is cheap), firms tend to appear greener (i.e. households have a higher estimate of greenness \hat{g}), although they are in fact less green (as seen in figure 2.4a). Figure 2.4c shows the evolution of capital. In “Perfect information”, firms remove brown capital very quickly to achieve greenness, leading to a noticeable drop in capital in the beginning. The

later rise in capital is driven by reallocating to green capital, which has a higher adjustment cost. This can be seen in figure 2.4d, where emissions initially drop rapidly in “Perfect information” and then stabilize despite capital rises, because the rise is mostly attributed to green capital. In models where households cannot observe greenness, both capital and emissions don’t follow sharp drops as in “Perfect information”. Although when greenwashing is mild (θ_E is high), emissions eventually drop to the same level as in “Perfect information”, the intertemporal difference in emissions suggests that information frictions are still costly. When greenwashing is strong (θ_E is low), emissions do not decline to the same level as in “Perfect information”, even after 10 years.

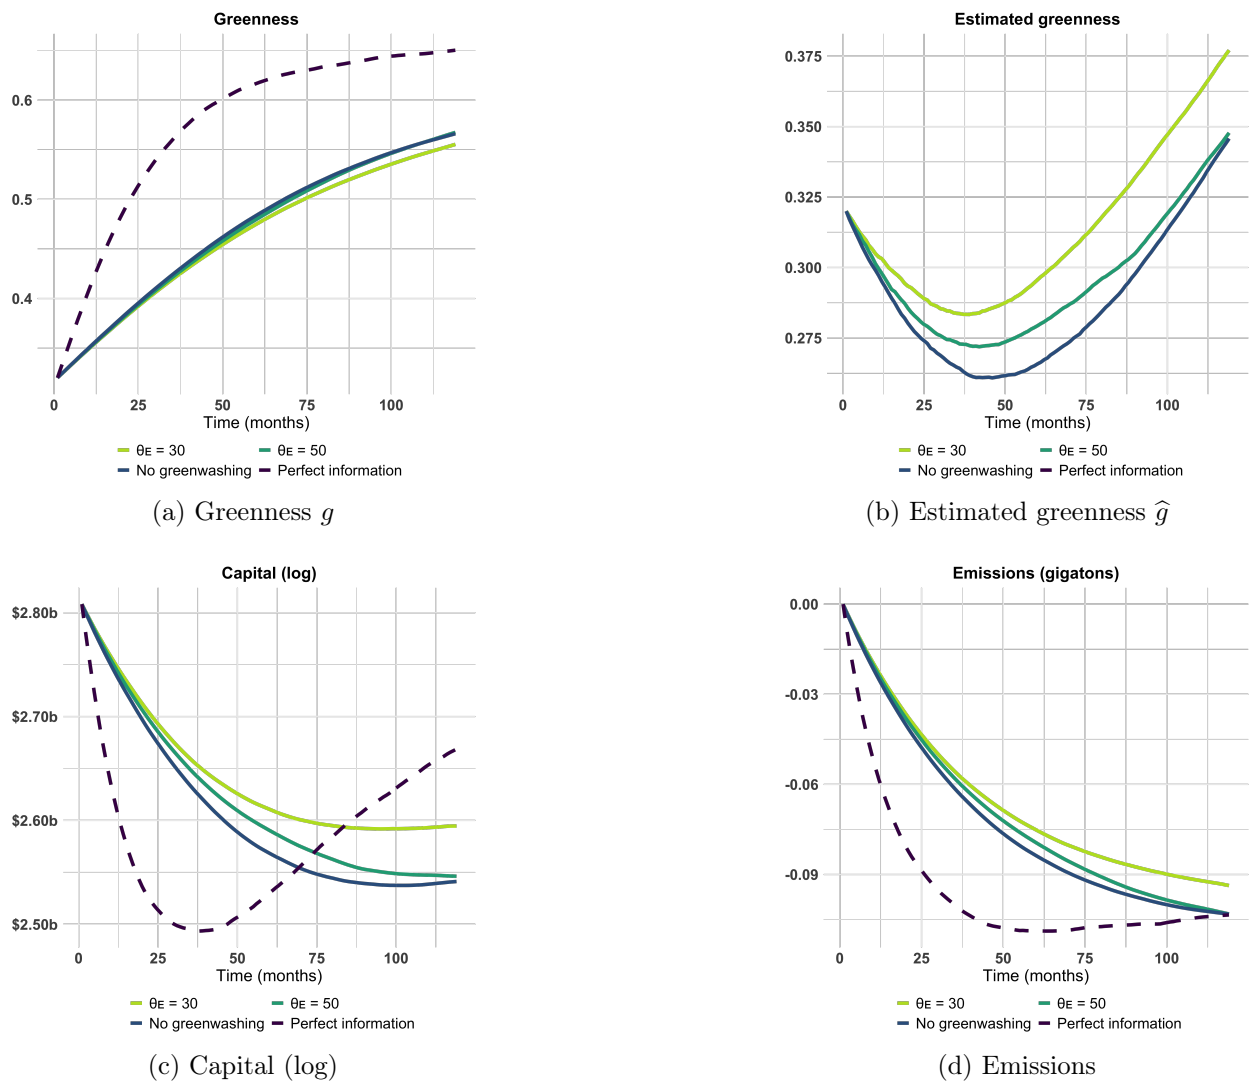


Figure 2.4: Variables of interest for different model specifications

In table 2.8, I summarize the results by comparing the outcomes at the end of the simulations (over 10 years). In all but one model, emissions converge to the same reduction, but firms take different approaches. In “Perfect information”, firms reallocate to more green capital, whereas in the other two cases, they hold less capital and thus emit less. When greenwashing is cheap, firms reduce emissions by -0.09 gigatons instead of -0.1 . Furthermore, in “Perfect information”, firms are “almost there” (reducing emissions by -0.08 gigatons) after 20 months, whereas in the model where greenwashing is cheap, it takes more than three times as long.

Table 2.8: Model Outcomes after 10 Years

	Perfect info.	$\theta_E = 30$	$\theta_E = 50$	No greenwashing
Greenness	0.65	0.56	0.57	0.57
Capital (log)	2.67	2.59	2.55	2.54
Emissions (gigatons)	-0.10	-0.09	-0.10	-0.10
Estimated greenness		0.38	0.35	0.35
Social cost of carbon \$bil	155	157	142	139
Time to -0.08 emissions (months)	20	70	61	55

In table 2.8, I also show the social cost of carbon. It is given by the present discounted value of scc_t (as defined in (2.32)) for ten years:

$$\mathbb{E} \left[\int_0^{10} \exp(-\rho t) scc_t dt \right]$$

The social cost of carbon is the lowest when greenwashing is not allowed (“No greenwashing”), and decreases in θ_E : the cost is \$139 billion when greenwashing is not allowed, whereas the cost jumps to \$157 billion when $\theta_E = 30$ (i.e. when greenwashing is cheap), implying a difference of \$18 billion. As a remark, when households can perfectly observe greenness (“Perfect information”), the social cost of carbon is \$155 billion, which is higher than the cost when greenwashing is not allowed (“No greenwashing”). This counter-intuitive result is due to the fact that the recapitalization of firms in “Perfect information” and zero spending on greenwashing lead to higher consumption (and thus lower marginal utility of consumption). As indicated in figure 2.4c, capital eventually rises in “Perfect information”.

To understand the interaction between firm characteristics, emissions, and greenwashing, I run the following regressions on the simulated data, similar to Fernández-Villaverde et al. (2024b):

$$\begin{aligned}
 \text{Percentile}(\widehat{g}_{i,t}) &= \beta_0 + \beta_1 \ln k_{i,t} + \beta_2 \theta_E + \gamma_t + \epsilon_{i,t} \\
 \Delta \text{Emission Intensity}_{i,t} &= \beta_0 + \beta_1 \text{Percentile}(\widehat{g}_{i,t}) + \beta_2 \theta_E + \gamma_t + \epsilon_{i,t} \\
 x_{E,i,t} &= \beta_0 + \beta_1 \ln k_{i,t} + \beta_2 \theta_E + \gamma_t + \epsilon_{i,t}
 \end{aligned} \tag{2.33}$$

where $\text{Percentile}(\widehat{g}_{i,t})$ is the percentile of the estimated greenness and $x_{E,i,t}$ is the greenwashing decision for firm i at time t . I report the results in table 2.9, which show that my model qualitatively replicates the empirical patterns found in section 2.3. In particular, larger firms tend to appear greener and greenwash more, whereas appearing greener doesn't imply a lower emission intensity. Furthermore, greenwashing penalty Θ does limit greenwashing, as the coefficient on θ_E is negative for the regression on x_E .

Table 2.9: Regression results from simulations

	Dependent Variables		
	Est. Greenness Percentile (1)	Δ Emission Intensity (2)	Greenwashing (3)
Capital (log)	74.984*** (0.090)		0.001*** (0.00000)
Greenwashing Penalty	0.100*** (0.005)	-0.0001*** (0.00000)	-0.00001*** (0.00000)
Est. Greenness Percentile		0.0002*** (0.00000)	
Time FE	Yes	Yes	Yes
Observations	192,000	192,000	192,000
R ²	0.782	0.573	0.355

Note: capital in \$ bil

*p<0.1; **p<0.05; ***p<0.01

2.7. Conclusion

This paper contributes to our understanding of the macroeconomic consequences of greenwashing by developing a novel general equilibrium model that incorporates distorted learning in a mean-field game setting. My analysis, motivated by empirical observations, reveals several important findings.

First, my model successfully replicates key empirical regularities observed in the data, including the tendency for larger firms to appear greener and the lack of correlation between environmental scores and future emission reductions. This alignment between model predictions and empirical facts lends credence to my theoretical framework.

Second, my quantitative analysis demonstrates that greenwashing can lead to significant economic inefficiencies and environmental consequences. Specifically, I find that when greenwashing is inexpensive (low θ_E), firms can take up to three times longer to reduce emissions compared to scenarios with perfect information. Even after a decade, economies with cheap greenwashing exhibit emissions that are 0.01 gigatons (0.1% of total US emissions) higher and are 9% less green than in scenarios where households can directly observe firms' environmental performance.

Third, my welfare analysis indicates that greenwashing can result in a loss of up to \$18 billion in social cost of carbon. This finding underscores the tangible economic costs associated with information asymmetries in environmental performance reporting.

Lastly, my paper naturally extends to other quantitative exercises. For example, one can adjust the carbon tax and green adjustment cost parameters, and analyze the transition paths. A higher carbon tax would increase the greennium and incentive more greenwashing, whereas low green adjustment cost may disincentive greenwashing. I leave these exercises to future research.

CHAPTER 3

A DEEP LEARNING ANALYSIS OF CLIMATE CHANGE, INNOVATION, AND UNCERTAINTY

This is joint work with Michael Barnett, William “Buz” Brock, Lars Peter Hansen, and Ruimeng Hu. My contribution is, in part, to the development of the economic model and computational methods, as well as the implementation of the latter.

3.1. Introduction

The characterization of optimal quantities, pricing outcomes, and social welfare implications under specified optimality criteria are the core focus of theoretical economic and financial market modeling. The outcomes from these modeling frameworks are often the emphasis for academic analysis and exploration, but such results can also be immensely consequential for government decision-making and policy action. As the complexity of modeling frameworks and policy questions increase, the limits of existing methodologies to provide sufficiently robust and accurate numerical solutions become ever more apparent. In this paper, we construct a deep learning algorithm by which we can solve economic and financial models that incorporate some of the most mathematically complex modeling elements. We focus on continuous-time, dynamic general equilibrium models with high-dimensional state spaces governed by jump-diffusion processes that are subject to model uncertainty concerns, model features that have proven to be particularly problematic for existing computational methods.

To contextualize our computational framework, we focus on a novel and particularly relevant climate-economics model setting. The potential consequences of climate change are becoming increasingly apparent, with the the Sixth Assessment Report (AR6) of the Intergovernmental Panel on Climate Change (IPCC) stating that “[i]t is unequivocal that human influence has warmed the atmosphere, ocean and land” and that anthropogenic climate change “has caused widespread adverse impacts and related losses and damages to nature and people, beyond natural climate variability.”²

²See Masson-Delmotte et al. (2021a) and Pörtner et al. (2022).

Given these concerns, policymakers and organizations are increasingly focusing on the necessity and feasibility of a transition to a carbon-neutral economy. Reports from the OECD, the IEA, the White House, McKinsey & Co., and Princeton University’s High Meadows Environmental Institute³, among many others, emphasize the need to heavily invest in cleaner production methods to reduce current carbon emissions, as well as the need to devote significant resources to R&D for developing new green technology to prevent future carbon emissions. These conclusions are based largely on scenario analysis aimed at achieving climate policy goals such as the 1.5 C° GMT temperature anomaly ceiling proposed by the 2015 Paris Climate Agreement. The implementation of a socially efficient carbon-neutral transition, taking into account the various frictions, risks, uncertainties, and endogenous feedbacks and responses related to climate change and climate policy action, is the key economic question that we address in this paper.

We develop and solve a dynamic general equilibrium framework with three types of capital: “dirty” capital that generates carbon emissions when used in production, “green” capital that produces no emissions but is initially less productive than dirty capital, and “knowledge” capital that increases the likelihood of a technology shock that augments the productivity of green capital. By focusing on capital stocks, our model incorporates dynamic features that play an important role in the transition to carbon-neutrality. First, the model framework leads to an emissions pathway that is “sticky” or persistent because dirty capital depreciates gradually and is costly to disinvest. Second, convex adjustment costs related to investment capital mean that the accumulation of new green capital can take significant time. Finally, the arrival of improved green technology depends on the stock of knowledge capital in the economy, which requires R&D investment across time and substitutes resources away from consumption, as well as from other types of investment.

Two important considerations in our analysis are the potential for thresholds or jumps in the climate damage and green innovation processes, as well as the substantial uncertainty about the central mechanisms in our model. We assume that the severity of economic damages from climate change is revealed via a stochastic jump process that is dependent upon the global mean temperature

³See OECD (2019), Bouckaert et al. (2021), White House (2021), McKinsey Global Initiative (2022), and Jenkins et al. (2021).

anomaly, and that the realization of a green technological innovation that enhances the productivity of the green capital occurs via a jump process that depends upon the level of knowledge capital accumulated through R&D investment. The uncertainty in our pertains to carbon-climate dynamics, or the mapping from carbon emissions into atmospheric carbon into temperature changes; the severity of the economic damage functions, or the negative impact on output due to changes in atmospheric temperature; and the technological innovation process, or the likelihood of a technology shock that augments the productivity of green capital via R&D investment. While each of these model components allows for risk in the form of stochastic realizations, we are interested in exploring alternative forms of uncertainty. In particular, we focus on the possibility that a given model is misspecified in a meaningful and unknown way. We explicitly incorporate these uncertainty considerations into the social planner’s decision problem by applying tools from dynamic decision theory. As a result, our analysis highlights the endogenous interdependencies and feedback effects that substantially impact optimal policy choices and social valuations coming from uncertainty aversion, with particular concerns related to shifts in climate damage severity and green technology.

Our general equilibrium framework with multiple endogenous state variables requires solving relatively high-dimensional PDEs with significant non-linearities due to the model uncertainty concerns and stochastic jump processes related to technological change and climate damages. As a result, our analysis requires computational methods that provide global solutions to accurately characterize the endogenous optimal policies and the social valuations of interest. Standard finite difference methods commonly used in economics and finance are not well equipped for such problems. We, therefore, develop and implement an algorithm using deep neural network methods, augmenting the deep Galerkin method-policy improvement algorithm setting (DGM-PIA) by incorporating key model parameters as pseudo-state variables. While using parameters as pseudo-state variables has been done in other settings⁴, to the best of our knowledge, we are the first to incorporate this feature into the DGM-PIA setting to improve the algorithm’s solution accuracy when confronting nonlinearities from jump processes and model uncertainty aversion. The use of pseudo-states is

⁴See Norets (2012); Duarte (2018), and Friedl et al. (2023) for examples of using pseudo-states in neural network frameworks related to structural model estimation and model sensitivity analysis.

critical for allowing the algorithm to accurately learn the monotonicity and other features related to the nonlinearity in our framework arising from model uncertainty aversion and stochastic jumps. As a result, the algorithm is able to handle multiple non-stationary, endogenous state variables with considerable non-linearity in a continuous-time, infinite horizon setting. The algorithm provides a toolbox that expands the ability of researchers to address questions not only in climate-economics, but more broadly in economics, finance, dynamic decision theory, and stochastic optimal control, where numerical solutions are necessary due to the “curse of dimensionality” and the complexities associated with heterogeneity, uncertainty, and nonlinearity in the model framework.

3.1.1. Deep Learning Literature

In recent years, deep learning algorithms, built on the neural network’s remarkable ability to represent and approximate high-dimensional functions and efficient gradient descent optimizers, have been very successful in many areas, ranging from computer vision and speech recognition to scientific computing (see, e.g. Carleo and Troyer (2017); Gao and Duan (2017); Han et al. (2018); Zhang et al. (2018); Han et al. (2019)). The climate change problem considered in this paper aims to study how choices for investment in clean and dirty capital, and R&D for technological innovation, are determined when facing uncertainties, which, mathematically, is a stochastic control problem. Along this direction, deep learning algorithms are roughly categorized into: direct parameterization, partial differential equations (PDEs), and forward-backward stochastic differential equations (FBSDEs) approach.

In this first category, the seminal work in high-dimension was proposed by Han and E (2016), which solves a global-in-time minimization problem in accordance with the utility of the control problem. Later, Bachouch et al. (2021) extended to the local-in-time approach combined with dynamic programming techniques. Han and Hu (2021) solve the control problem with aftereffects, modeled by stochastic differential delayed equations, using the same spirit of Han and E (2016) but with advanced network architectures, and Carmona and Laurière (2022) extended to mean-field control problems.

In the PDE approach, the stochastic control problem will first be reformulated into a Hamilton-

Jacobi-Bellman (HJB) PDE using the dynamic programming principle, and then solved by deep learning algorithms. For generic PDEs, Sirignano and Spiliopoulos (2018) proposed the deep Galerkin method (DGM), and Raissi et al. (2019) proposed the physics-informed neural networks (PINNs), roughly at the same time. Both approximate solutions to PDEs by training neural networks to minimize the residuals coming from the initial conditions, the boundary conditions, and the PDE operators. Saporito and Zhang (2021) followed this idea and solved path-dependent PDEs using recurrent neural networks. Later, Al-Arabi et al. (2022) and Duarte et al. (2023) extended the DGM to deal with HJB equations in their unsimplified primal form and solved for the value function and the optimal control simultaneously by characterizing both as deep neural networks.

For semi-linear parabolic PDEs, which correspond to stochastic control problems with uncontrolled volatility, E et al. (2017) and Han et al. (2018) proposed a deep BSDE solver which, to the best of our knowledge, is the first work in this area and has inspired much follow-up work. They recast the solution of a semi-linear parabolic PDE into a global-in-time optimization problem based on the associated BSDEs via the non-linear Feynman Kac formula and variational form. Huré et al. (2020) dealt with the same associated BSDE but obtained the resolution by solving backward in time using a sequence of small optimization problems and termed it as deep learning backward dynamic programming (DBDP). Another work focusing on semi-linear PDE is the deep splitting method proposed by Beck et al. (2021), where the partial differential operators are split into the linear part and the nonlinear part, with the nonlinear part propagating first by freezing the solution and the linear part then being taken care of by an approximate Feynman-Kac representation. When the control appears in the state dynamics' diffusion coefficient, this leads to a fully nonlinear PDE. In this direction, in the same spirit of E et al. (2017) and Han et al. (2018), Beck et al. (2019) proposed a DL algorithm by solving the corresponding second-order BSDE, and Pham et al. (2021) extended the DBDP idea in Huré et al. (2020) by further approximating the Hessian matrix using auto differentiation of first derivatives' neural networks.

In the FBSDE approach, the control problem will first be reformulated into a coupled forward-backward system using the stochastic Pontryagin maximum principle. The system is, in general,

hard to solve numerically, let alone in high dimensions, due to its coupled nature and opposing directions: one with an initial condition running forward in time and one with a terminal condition running backward in time. Han and Long (2020) solved the fully-coupled FBSDE by extending the deep BSDE solver with convergence analysis subject to neural networks' universal approximation, and Ji et al. (2020) further extended the results with the Picard iteration method.

Recently, attention has also been drawn to stochastic control problems with jumps. The problem is more involved as the PDE becomes partial integro-differential, thus non-local; and the FBSDE system now contains a Lévy process. For such problems, Boussange et al. (2022) solved the associated non-local PDE by extending the deep splitting method in Beck et al. (2021) and multilevel Picard approximation method in Hutzenthaler et al. (2019); Castro (2022) solved the associated forward-backward system by extending the DBDP method in Huré et al. (2020); and Gnoatto et al. (2022) solved the same system in the same spirit of the deep BSDE solver, proposed in E et al. (2017) and Han et al. (2018). These settings all focus on the semi-linear problem, where the control only appears in the drift coefficient. Very recently, Lu et al. (2024) proposed an actor-critic reinforcement learning algorithm to address the fully nonlinear control problem with jumps.

When using neural networks to parameterize the quantity of interest, usually the value function or the control policy, people may want to incorporate domain knowledge, such as guaranteed monotonicity or convexity of the learned function with respect to some of the input variables. There have been fruitful studies on how to build network structures in order to fulfill such requirements, see for instance, Sill (1997), Gupta et al. (2016), Wehenkel and Louppe (2019) and Runje and Shankaranarayana (2022).

3.1.2. Climate Economics Literature

Our paper builds on and contributes to a number of important areas of research across economics, finance, geoscience, and applied mathematics. The implications of anthropogenic emissions have been a central focus of geoscientists for many decades, beginning with the seminal work of Arrhenius (1896). Recent work has focused on characterizing the dynamic relationship between carbon emissions, atmospheric carbon concentration, and temperature change via pulse experiments (Joos et al.,

2013; Geoffroy et al., 2013; Eby et al., 2009), deriving simplified emulators or approximations of these complex relationships for the use of policymakers (Matthews et al., 2009; Pierrehumbert, 2014; MacDougall et al., 2017), as well as quantifying the dynamic and stochastic features of carbon-climate dynamics (Ricke and Caldeira, 2014; Palmer and Stevens, 2019). These components are critical inputs into our model framework and uncertainty analysis.

The climate economics literature has given substantial focus to estimating the economic consequences of climate change and deriving a Social Cost of Carbon (SCC), or the cost to social welfare of emitting an additional amount of carbon emissions into the atmosphere. Economists have estimated the economic costs generated by observed climate change for numerous economic sectors, regions, and dimensions of the economy (Dell et al., 2012; Hsiang et al., 2017; Colacito et al., 2019; Allen et al., 2019; Carleton and Greenstone, 2021). Theoretical modeling and integrated assessment model analysis to derive SCC valuations have considered for numerous economic mechanisms (Golosov et al., 2014; Acemoglu et al., 2016; Nordhaus, 2018; Cai and Lontzek, 2019,?) and proposed climate damage function approximations and frameworks (Lenton et al., 2008; Weitzman, 2012; Cai et al., 2015; Drijfhout et al., 2015; Ritchie et al., 2021).

Recent work has begun to examine important dimensions related to interconnected climate change and economic model uncertainty (Olson et al., 2012; Lemoine and Traeger, 2014; Hassler et al., 2018; Nordhaus, 2018; Dietz and Venmans, 2019; Barnett et al., 2020; Rudik, 2020; Barnett et al., 2022; Barnett, 2023). Importantly, many of these studies of uncertainty have exploited the powerful toolset developed in dynamic decision theory (Anderson et al. (2003); Maccheroni et al. (2006); Hansen and Sargent (2007); Klibanoff et al. (2009); Hansen and Miao (2018); Barnett et al. (2020); Cerreia-Vioglio et al. (2021)), allowing researchers to account for model uncertainty explicitly in the decision-maker's problem within the model.

This paper also links to an important area of macroeconomic research. This includes the foundational work on endogenous economic growth (Brock and Mirman, 1972; Baumol, 1986; Lucas Jr, 1988; Romer, 1990), as well as recent analysis of the social and private benefits of innovation (Bloom et al., 2019; Lucking et al., 2019), and considerations for the transition to a green econ-

omy (Acemoglu et al., 2012, 2016; Jaakkola and Van der Ploeg, 2019). In addition, the connection between macroeconomics and asset pricing in production-based asset pricing (Brock, 1982; Cochrane, 1991; Jermann, 1998), as well as the importance links to economic growth and innovation (Papanikolaou, 2011; Kogan and Papanikolaou, 2014; Kung and Schmid, 2015) have important implications for the social valuations we derive in our climate-economics-innovation linked framework.

3.2. A Deep Learning Algorithm for Nonlinear Problems

We start by outlining the Hamilton-Jacobi-Bellman equation used to characterize the types of dynamic, recursive optimization problems we are interested in solving. We then detail the implications that jump processes and model uncertainty aversion have for our baseline HJB equation. With these pieces in place, we discuss the implications of high-dimensionality in the state space, further motivating the numerical algorithm needed. We then detail the algorithm we construct to solve these types of models, accounting not only for the nonlinearities but also the high dimensionality. In particular, we show how including “pseudo”-state variables significantly improves the accuracy of global solutions in such complex settings.

3.2.1. Standard HJB equation

For simplicity and completeness, we begin by outlining the set-up for a standard dynamic, recursive optimization problem. We consider a social planner’s problem where the objective is to maximize lifetime, discounted expected utility in continuous time over an infinite horizon. Following the relatively standard heuristic derivation, omitted for conciseness, the social planner’s problem can be written as an autonomous partial differential equation (PDE) known as a Hamilton-Jacobi-Bellman (HJB) equation, which takes the following form:

$$0 = -\delta V(x) + \sup_{\alpha} \left\{ U(x, \alpha) + \frac{\partial V}{\partial x'}(x) \mu(x, \alpha) + \frac{1}{2} \text{trace} \left[\sigma(x, \alpha)' \frac{\partial^2 V}{\partial x \partial x'}(x) \sigma(x, \alpha) \right] \right\}.$$

where $V(x)$ represents the value function capturing the optimized social welfare of the social planner, x and α are realizations of a vector of relevant state variables X_t , whose evolution is given by

stochastic diffusion process with drift $\mu(x, \alpha)$ and volatility $\sigma(x, \alpha)$, and control choices A_t , δ is the subjective discount rate, $U(X_t, A_t)$ is the flow utility function. Various computational algorithms have been devised to derive viscosity solutions to such relatively standard PDEs when analytical solutions are not available. Our focus is on going beyond such settings as we outline below.

3.2.2. Nonlinearity from Jumps and Model Uncertainty

We are interested in augmenting the basic setting outlined above in ways that are central for climate-economics and climate-finance model frameworks, as well as other stochastic optimal control settings in economics, finance, and applied mathematics. The additional model components we focus on are jump processes and model uncertainty aversion, which introduce endogenous non-linearities that often rule out analytic solutions and many existing numerical methods. We outline the impact that these model features have below.

Stochastic Jump Processes

We begin with the introduction of stochastic jumps for the state vector X_t . We assume in our setting that there are ℓ discrete realizations of the jump process from a set \mathcal{L} , with a jump intensity function $\mathcal{J}^\ell(x) = \pi^\ell \mathcal{J}(x)$. Note π^ℓ is the prior probability of jump realization ℓ occurring if a jump takes place and $\mathcal{J}(x)$ is the state-dependent jump arrival rate. Allowing for jumps in our setting leads to the addition of the following term to our HJB equation

$$\sum_{\ell=1}^{\mathcal{L}} \mathcal{J}^\ell(x) \left[V^\ell(x) - V(x) \right].$$

where $V^\ell(x)$ denotes the post-jump value function for jump realization ℓ . Depending on the function form of the arrival rate and jump realizations, the jumps can add an additional endogenous, convexly nonlinear element to our HJB equation that complicates the ability to derive numerical solutions to our problem.

Brownian Misspecification

Next, we outline the impact of Brownian model misspecification. Following from Girsanov theory, diffusion misspecification comes in the form of a drift distortion related to a positive martingale

with expectation equal to one which parameterizes changes in the probability measure. With this, we use robust preferences of the form explored by Hansen and Sargent (2011) which augments the planners preferences and optimization problem to include a minimization choice based on a quadratic penalty term over h coming from a local measure of relative entropy or Kullback-Leibler divergence to capture “model uncertainty aversion”. The problem augments our baseline HJB equation as follows

$$\min_h \frac{\partial V}{\partial x'}(x) [\mu(x, \alpha) + \sigma(x, \alpha)h] + \frac{1}{2} \text{trace} \left[\sigma(x, \alpha)' \frac{\partial^2 V}{\partial x \partial x'}(x) \sigma(x, \alpha) \right] + \frac{\xi}{2} h' h$$

where ξ is a scalar parameter capturing “model uncertainty aversion” and will be determined as in a robust Bayesian setting by considering implied ex-post probabilistic distortions generated by the choice of ξ . Optimizing over h gives the optimal control

$$h^* = -\frac{1}{\xi} \sigma(x, \alpha)' \frac{\partial V}{\partial x}(x).$$

As with jumps, Brownian misspecification can add an additional endogenous nonlinear element to our HJB equation, depending on assumed functional forms, which also complicates the ability to derive numerical solutions to our problem.

Jump Misspecification

Finally, we highlight the impact of jump model misspecification. To incorporate this type of misspecification, we again follow the robust preferences structure of Hansen and Sargent (2011). We therefore introduce non-negative functions g^ℓ which alter the jump intensity to be of the form $\mathcal{J}^\ell(x)g^\ell(x)$. The planner’s preferences and optimization problem are again augmented to include a minimization choice based on a local measure of relative entropy or Kullback-Leibler divergence over the $g^\ell(x)$ to capture “model uncertainty aversion” of this form. The problem augments our baseline HJB equation as follows

$$\min_{g^\ell \geq 0} \sum_{\ell=1}^L \mathcal{J}^\ell(x) g^\ell(x) \left[V^\ell(x) - V(x) \right] + \xi \sum_{\ell=1}^L \mathcal{J}^\ell(x) \left[1 - g^\ell(x) + g^\ell(x) \log g^\ell(x) \right]$$

where ξ is our aforementioned a scalar parameter capturing “model uncertainty aversion”. Optimizing over the $g^\ell(x)$ gives the optimal control

$$g^{\ell*}(x) = \exp\left(-\frac{1}{\xi} \left[V^\ell(x) - V(x)\right]\right).$$

Again, as in our previous two modifications, jump misspecification can add an additional endogenous nonlinear element to our HJB equation, depending on assumed functional forms. Moreover, this nonlinear interacts with the existing nonlinear adjustment coming from the introduction of jumps even without misspecification concerns. Together, this additional component further complicates the ability to derive numerical solutions to our problem.

3.2.3. Augmented HJB equation

Our climate-economics setting incorporates each of these elements simultaneous, with each playing important parts of the considerations for economic damages from climate change, technological innovation and the carbon-neutral transition, and the model uncertainty associated with these elements, as well as the broader climate-economics dynamics underlying the modeling framework. Taken together, and plugging in the model uncertainty-associated minimizers, our baseline HJB equation is augmented to take the following form

$$0 = -\delta V(x) + \sup_{\alpha} \left\{ U(x, \alpha) + \frac{\partial V}{\partial x'}(x) \mu(x, \alpha) + \frac{1}{2} \text{trace} \left[\sigma(x, \alpha)' \frac{\partial^2 V}{\partial x \partial x'}(x) \sigma(x, \alpha) \right] - \frac{1}{2\xi} \frac{\partial V}{\partial x'}(x) \sigma(x, \alpha) \sigma(x, \alpha)' \frac{\partial V}{\partial x}(x) \right\} + \xi \sum_{\ell}^L \mathcal{J}^\ell(x) \left[1 - \exp\left(-\frac{1}{\xi} \left[V^\ell(x) - V(x)\right]\right) \right]$$

where we have included jumps and model misspecification of both forms, having plugged in the optimized distortions h and $g^\ell(x)$. The challenge we face is constructing an algorithm that can accurately capture the nonlinearities resulting from each of these model features, as well as their interactions, in a dynamic climate-economics model setting.

3.2.4. High Dimensionality

The final component we must consider for our computational algorithm is confronting the “curse of dimensionality”, or that the complexity of deriving numerical solutions exponentially increases in the number of dimensions. The integrated structure of theoretical climate-economics and finance models requires a high dimensional state space in order to rigorously and accurately account for relevant model features. As noted by Han et al. (2018), numerical solutions using standard finite difference or finite element methods are almost always unavailable for problems where the number of dimensions is greater than or equal to four⁵. Incorporating model uncertainty and jumps, which introduces endogenous non-linearities as noted before, can further exacerbate the computational burden in high dimensions.

3.2.5. Implementing Neural Nets for Numerical Solutions

We therefore propose a novel implementation of deep neural networks for deriving the numerical solutions to HJB equations by focusing on an algorithm designed not only to deal with high-dimensionality, but also the nonlinearities and complexities associated with jumps and model uncertainty as outlined before. Only very recently have deep learning and neural network methods been applied to solving dynamic economic models. Because of the ability of these methods to provide global solutions for problems where other computational solutions methods become infeasible or simply fail, such as with high dimensional state spaces or significant model nonlinearities, deep learning solution methods are seen as a potentially “game-changing” toolset. By approximating the representative agent’s value function and optimal controls with deep neural networks, the “curse of dimensionality” can be ameliorated due to the representation of functions in a compositional form, rather than by the standard additive form resulting from finite difference and element methods. Applications in macroeconomics (Fernandez-Villaverde et al., 2020; Maliar et al., 2021; Azinovic et al., 2022; Payne et al., 2024) and finance (Duarte et al., 2023; Sauzet, 2021) implementing these types of solutions methods have begun to demonstrate the potential value and importance for studying key problems in the literature more broadly.

⁵There are various alternative approaches that have been explored for solving problems with more than 4 state variables, of which deep learning algorithms are a recent one that appears to have substantial promise.

We now outline our algorithm, and provide pseudo-code and further details about implementation in the appendix. We build upon the deep Galerkin method-policy improvement algorithms (DGM-PIA) proposed in Al-Arabi et al. (2022) and Duarte et al. (2023) to solve the aforementioned HJB equations. We first illustrate the algorithm on a generic HJB equation for $V(\mathbf{x})$:

$$-\delta V(\mathbf{x}) + \sup_{\boldsymbol{\alpha} \in A} \{L^{\boldsymbol{\alpha}} V(\mathbf{x}) + f(\mathbf{x}, \boldsymbol{\alpha})\} = 0,$$

where \mathbf{x} and $\boldsymbol{\alpha}$ denote the state and control variables, A is the control space, the differential operator $L^{\boldsymbol{\alpha}}$ is the infinitesimal generator of the controlled state process $\mathbf{X}^{\boldsymbol{\alpha}}$, f is the utility function and δ is the discount factor. DGM-PIA solves for the value function V and the optimal control $\boldsymbol{\alpha}$ simultaneously by parameterizing both as deep neural networks V^{θ} and $\boldsymbol{\alpha}^{\varphi}$. Then, the networks are trained by taking alternating stochastic gradient descent steps for the two functions. Let $\boldsymbol{\alpha}^{\varphi_0}$ (as a function of \mathbf{x}) be the initial control parameterized by the neural net, at stage n , the algorithm contains two steps:

Step 1. Find a solution to the linear PDE

$$-\delta V^{\theta_n}(\mathbf{x}) + L^{\boldsymbol{\alpha}^{\varphi_n}} V^{\theta_n}(\mathbf{x}) + f(\mathbf{x}, \boldsymbol{\alpha}^{\varphi_n}(\mathbf{x})) = 0,$$

for the fixed control $\boldsymbol{\alpha}^{\varphi_n}$, by updating θ_n via minimizing

$$L_V(\theta) = \|\delta V^{\theta}(\mathbf{x}) + L^{\boldsymbol{\alpha}^{\varphi_n}} V^{\theta}(\mathbf{x}) + f(\mathbf{x}, \boldsymbol{\alpha}^{\varphi_n}(\mathbf{x}))\|^2.$$

Step 2. Update the policy corresponding to

$$\boldsymbol{\alpha}^{\varphi_{n+1}}(\mathbf{x}) \in_{\boldsymbol{\alpha} \in A} \{L^{\boldsymbol{\alpha}} V^{\theta_n}(\mathbf{x}) + f(\mathbf{x}, \boldsymbol{\alpha})\},$$

for the fixed value function V^{θ_n} , by update φ_{n+1} via minimizing

$$L_{\alpha}(\varphi) = - \int_{\Omega} \left[L^{\alpha^{\varphi}} V^{\theta_n}(\mathbf{x}) + f(\mathbf{x}, \alpha^{\varphi}(\mathbf{x})) \right] \nu(\mathbf{x}),$$

where $\nu(\mathbf{x})$ is a probability measure on the domain Ω of \mathbf{x} characterizing the different regions' relative importance.

Augmenting the Algorithm with “Pseudo”-State Variables

To effectively account for the nonlinearities resulting from model uncertainty and jumps, we extend our algorithm by incorporating “pseudo”-state variables that are used as inputs for the parameterized neural network. Specifically, we alter the stochastic gradient descent step in the DGM-PIA algorithm to randomly sample points from the pseudo-state space, defined in our setting by values of the model uncertainty aversion parameter ξ and the potential realizations of the jump processes denoted by $\ell \in L$, as well as the standard state space. We are therefore searching for solutions over a range of ξ and ℓ values, even though there is no dynamic evolution associated with these pseudo-states. This enhances the algorithm’s ability to characterize the nonlinearities associated with the jump and model uncertainty model mechanisms, allowing us to find accurate global solutions to our problem, in particular, those aligned with economic intuition.

The reason why the use of pseudo-state variables improves the algorithm’s accuracy is actually fairly intuitive. By sampling from the pseudo-state space, we provide the neural networks access to relevant information over which to train that would not be available if we were to train individual networks for each model solution across different sets of parameters. Importantly, given our choice of pseudo-state variables, this additional cross-sectional information precisely targets the model components that are the most complex and difficult to characterize. Because the algorithm samples from the pseudo-state space, the network tries to learn the functional relationship with respect to the pseudo-state variables, similar to an operator learning framework. A tendency of neural networks is to exhibit implicit regularization, meaning they tend to search for simpler function solutions that fit the data well, such as those with monotonicity, even without imposing such constraints. If the

neural networks finds a form of regularity consistent with the data, it will likely prefer this solution to other alternatives, implicitly providing additional structure to the overall solution. With this additional structure over the pseudo-state space, as the the neural network updates the solution for one set of parameter values, it will impact the solution for other parameter values. These updates are generally beneficial because the model solutions we are searching for are usually continuous and smooth in the state and parameter space. As a result, the solution for one set of parameters is likely to be relatively close to the solution for a slight perturbation to the parameter values. This additional information and structure provided by the inclusion of pseudo-states appears to be particularly valuable in allowing the neural networks to more accurately identify and fit the types of complex nonlinearities encompassed by our setting.

Our pseudo-state approach extends the deep Galerkin method, which aims to solve PDEs while maintaining a monotonic relationship with certain model parameters. This fundamentally differs from neural operator learning methods like DeepONet, which learns the functional relationship between model parameters and the solution and can be considered a form of operator learning (Kovachki et al., 2024)⁶. Rather than learning an operator from a limited set of parameters and corresponding PDE solutions, our approach samples a broad range of parameters and solves the corresponding PDEs directly. This approach emphasizes accurately solving the PDEs across a range of parameters while ensuring certain monotonicity properties, thereby distinguishing it from the operator learning framework.

3.2.6. The Contribution in Context

To be concrete, we are unaware of any existing contributions in literature related to deep learning algorithms on solving dynamic optimal control problems with jumps and model uncertainty. Moreover, while there are non-deep learning algorithms that can handle problems of dimension 4 or higher, they are often tailored for specific problems, taking advantage of particular model structures. In contrast, our model solution setting is very general and complex, and thus a contribution of ours is to develop a “toolbox” for highly nonlinear control problems with jumps and model uncertainty.

⁶Li et al. (2020) and Duarte et al. (2023) provide additional examples of operator learning.

Though existing work has explored settings with (potentially many) more state variables than ours, a number of settings require assumptions such as state variables that are symmetric, stationary, or exogenously specified, or linearity in the model framework, to maintain tractability with such scale. And while some frameworks can handle non-trivial nonlinearities or forms of uncertainty analysis, to the best of our knowledge ours is the first using neural network solution method to confront the endogenous nonlinearities resulting from the application of dynamic decision theory to model uncertainty aversion, combined with stochastic jump processes. Developing a methodology that handle such complexities is essential to analyzing the type of climate-economics model with model uncertainty aversion broadly conceived that we put forward next.⁷

We address these issues in our framework by implementing an extended deep Galerkin method algorithm that considers critical parameters as additional “pseudo” state variables in the network input, an extension of the DGM-PIA setting not previously explored in the literature to our knowledge. Because of this, we are able to derive global numerical solutions for a continuous-time, infinite horizon setting with significant non-linearities, multiple endogenous state variables, and stochastic jumps while incorporating aversion to model uncertainty via the toolset of dynamic decision theory. This is critical to our analysis, as exploring transition dynamics for endogenous optimal policy responses is at the heart of understanding and analyzing outcomes related to climate change, model uncertainty, and the transition to carbon-neutrality. Furthermore, our algorithm can still handle high-dimensional PDEs, attenuating the “curse of dimensionality” in such cases. As such, the resiliency of our algorithm to various modeling complexities, while still being able to scale to higher-dimensional settings, opens the door to exploring models with regional, household, and firm heterogeneities across technologies, economic frictions, policy objectives, and other sources. Thus, our numerical algorithm not only enriches our ability to explore a more extensive set of questions

⁷For example, while an improvised finite difference method using neural networks to solve linear systems may help mitigate the curse of dimensionality, it does not produce satisfactory results for complex models with strong nonlinearities as in our setting. Also, due to the strong nonlinearities in our model, naively parameterizing the value function using a neural network and minimizing the loss associated with the PDE operator (like the standard deep Galerkin method) will not produce accurate enough results. While an extended version of the deep Galerkin method (Al-Arabi et al., 2022; Duarte et al., 2023) is indeed helpful in improving accuracy, it does not effectively preserve the value function’s monotonicity with respect to certain parameters which are supported by economic arguments, nor does it sufficiently handle the nonlinearities arising from our application of dynamic decision theory to address concerns about model uncertainty.

and models in climate-economics and climate finance than before, specifically those that apply dynamic decision theory to address model uncertainty, but also more broadly in economics, finance, and other problems applying dynamic stochastic optimal control.

3.3. Climate-Economics Model Example

We now outline the example economy used to evaluate our numerical solution method. The model incorporates components related to climate change dynamics, economic damages from climate change, economic production technologies, technological innovation through R&D investment, and preferences that include model uncertainty aversion. We outline each of these model pieces in what follows, and then derive the central theoretical results before presenting the numerical solutions to the model derive using our computational algorithm.

3.3.1. Climate Dynamics

We begin with relatively simple climate dynamics based on the simplified approximations of long-run climate change outcomes developed by Matthews et al. (2009) and validated by Friedlingstein (2015). Specifically, recent geoscientific analysis has shown that the output of large-scale atmosphere-ocean general circulation models (AOGCM) can be approximated by a proportional relationship between temperature anomaly and cumulative carbon emissions:

$$\text{temperature anomaly} \approx \underline{\text{climate sensitivity}} (\theta) \times \text{cumulative emissions.}$$

For our analysis, we use a stochastic variant of this relationship given as follows:

$$dY_t = \mathcal{E}_t (\bar{\theta} dt + \varsigma dW_t^y),$$

where Y_t is the global mean temperature anomaly, in degrees Celsius, relative to the preindustrial level, \mathcal{E}_t is global carbon emissions, W_t is a vector Brownian motion process with filtration \mathcal{F}_t and scalar components W_t^y , W_t^d , W_t^g and W_t^r , ς is the volatility loading for temperature, and $\bar{\theta}$ is the average of the possible climate sensitivity proportionality constants $\theta(m)$ from a set of possible values denoted by $m = 1, \dots, M$, with equal-weighted likelihood given to each model. Importantly,




Figure 3.1: Implied TCRE coefficients across time and models, capturing carbon and temperature uncertainty. The top figure shows the percentiles for temperature responses to emission pulses for all carbon and temperature model combinations. The bottom figure is a histogram for the exponentially weighted average responses of temperature to an emissions pulse from 144 different models using a rate $\delta = .01$.

this long-run approximation implies a permanent impact of carbon emissions⁸. In addition, concerns about model uncertainty naturally arise in this setting given the variation in estimated $\theta(m)$ values, known as the transient climate response to cumulative emissions (TCRE) parameter, found across

⁸A reasonable approximation in our view given the estimated lifespan of atmospheric carbon emissions found by Pierrehumbert (2014) and others.

large-scale AOGCMs⁹.

In addition, the inclusion of Brownian shocks serves a few purposes. First, it is consistent with the seminal work of Hasselmann (1976)¹⁰, which started the field of stochastic climate modeling to characterize the fast/slow time scale separation of climate variability, and the more recent work of Palmer and Stevens (2019) as a means of improving predictability and approximating variation coming from chaotic models of climate dynamics. Second, Brownian shocks can be seen as a crude way of representing the unmodelled “complex interplay of multiple physical and biogeochemical processes” underlying the TCRE proportionality result that is “not amenable to a simple physical explanation” as recently shown by Gillett (2023). Third, the use of Brownian shocks means that the true climate model is “disguised” in the statistical sense that a small set of observations do not immediately reveal the true model, allowing for a meaningful uncertainty analysis.

While the proportionality relationship used here is based on a relationship of carbon emissions to temperature change that is born out at a time horizon of 10 years and beyond, Dietz and Venmans (2019); Barnett et al. (2022); Barnett (2023) and others have shown that this relationship is well suited for frequencies as short as one year. We follow Barnett et al. (2022) and Barnett et al. (2023b) and use pulse experiment results of Joos et al. (2013) and Geoffroy et al. (2013) to build the set of climate sensitivity models for our analysis.¹¹ Figure 3.1 shows the dynamic pathways and the implied TCRE parameters for the 144 models produced from this exercise. We can see from the pathways that the temperature response peaks around 10 years, and flattens out thereafter. The histogram implied TCRE parameters, also shown in Figure 3.1, are quite dispersed, with values ranging from around 1 to just below 3, with the average value being 1.86.




Figure 3.2: Range of possible damage functions for two cases with different jump thresholds. The solid lines show the average values, and the shaded regions give the range of possible values for $\exp(-n)$, which measures the proportional reduction of the productive capacity of the economy. The blue line and region show the damage function curvature when the jump occurs at $Y_t = \bar{y} = 2.0$. The red line and region show the damage function curvature when the jump occurs at $Y_t = \underline{y} = 1.5$. The black dashed lines indicate the values of Y_t for the upper and lower jump thresholds for the temperature anomalies.

3.3.2. Climate Damages

Our model of economic damages from climate change, or climate damages, follow the piece-wise log quadratic specification developed in Barnett et al. (2023b). Damages N_t are assumed to impact capital, output, investments, and consumption in a proportional manner, and allows for a change in the evolution of damages that is triggered by a jump process realization. Uncertainty about climate damages is connected to the values of \hat{Y} , which is defined by $\hat{Y} = Y + \bar{y} - \tilde{Y}$ for $\underline{y} < \tilde{y} \leq \bar{y}$, and

⁹Matthews et al. (2009), Friedlingstein (2015), and Masson-Delmotte et al. (2021b) all provide insightful representations and details related to this climate sensitivity uncertainty result.

¹⁰Franzke et al. (2022) surveys Hasselman’s work, as well as the work of Syukuro Manabe and Giorgi Parisi, who were jointly awarded the 2021 Nobel Prize in Physics for work on modeling complex physical systems.

¹¹Joos et al. (2013) examines carbon dynamics variation and uncertainty based on responses of atmospheric carbon concentration to emission pulses for several alternative Earth System models. Geoffroy et al. (2013) provides temperature dynamics variation and uncertainty based on approximate dynamics relating the log of atmospheric carbon to future temperature, following the Arrhenius equation. We take 9 atmospheric carbon pulse responses as inputs into 16 temperature dynamics approximations to derive our set of carbon-climate TCRE models.

$\lambda_3(\ell)$, for which there are alternative values denoted by $\ell = 1, \dots, L - S$. There is a stochastic jump process that determines dynamically which value of $\lambda_3(\ell)$ will be realized for the damage function and at the value of \tilde{Y} the jump will be realized. Specifically, the jump process has $L - S$ absorbing states, where each state ℓ corresponds to a value of $\lambda_3(\ell)$. Before the realization of the jump, each $\lambda_3(\ell)$ value has an equal-weighted likelihood of occurring. When the jump is realized, the value of $\lambda_3(\ell)$ is revealed and the true damage function curvature becomes known to the social planner in the model. The likelihood of a damage function jump is determined by a jump intensity $\mathcal{J}_n^\ell(y)$ that is increasing in the endogenous temperature anomaly state variable y and is localized to the range of temperature values $[\underline{y}, \bar{y}]$. At the time of the jump τ , the state variable \hat{Y} shifts from the temperature anomaly $Y = \tilde{Y}_\tau$ to the damage jump threshold \bar{y} . Thus \hat{Y} is a state variable that shifts from capturing just the temperature anomaly to a jump-adjusted measure capturing the derivative shift of the damage function as well. The uncertainty is associated both with the probability distribution on possible $\lambda_3(\ell)$ model realizations, as well as the probability of when a damage jump will occur in the $[\underline{y}, \bar{y}]$. We leave the full functional form for the Appendix, but the evolution of log damages is

$$d \log N_t = \begin{cases} \left(\lambda_1 + \lambda_2 \hat{Y}_t \right) \mathcal{E}_t \left(\bar{\theta} dt + \varsigma dW_t^y \right) + \frac{\lambda_2 |\varsigma|^2 (\mathcal{E}_t)^2}{2} dt, & t \leq \tau, \\ \left(\lambda_1 + \lambda_2 \hat{Y}_t + \lambda_3(\ell) \left(\hat{Y}_t - \bar{y} \right) \right) \mathcal{E}_t \left(\bar{\theta} dt + \varsigma dW_t^y \right) + \frac{(\lambda_2 + \lambda_3(\ell)) |\varsigma|^2 (\mathcal{E}_t)^2}{2} dt, & t \leq \tau. \end{cases}$$

For our main setting, we assume $[\underline{y}, \bar{y}] = [1.5, 2.0]$, consistent with the critical temperature thresholds used by the IPCC and the 2015 Paris Climate Agreement. Figure 3.2 shows the implied climate damages across all $\lambda_3(\ell)$ values for the jump realization occurring at the temperature anomaly values of $\tilde{y} = \underline{y} = 1.5$ and $\tilde{y} = \bar{y} = 2$ which bound our range for the damage jumps occurring. The red shaded region shows the range of possible damage outcomes across the possible values of $\lambda_3(\ell)$, and the black line gives the mean value. While there is significant discussion about the likelihood of tipping points or thresholds at the global scale captured by discrete jumps or shifts in climate damages¹², our setting is instead a smooth shift where the planner realizes whether the damage function is more convex than previously known. In this sense, we view our framework

¹²See Brook et al. (2013) and Levitan (2013) for examples on this discussion.

as a valuable tool for characterizing the possibility and uncertainty of potentially severe, nonlinear climate damage outcomes that are highly relevant for the optimal policy decisions of a social planner confronting climate change. Moreover, the structure of our uncertainty is novel with regards to much of the existing climate-economics literature, as the information dynamics in most settings is static, or uncertainty about the damage function is never resolved.

3.3.3. Production

For the economic component of the model, we assume there are two sectors producing perfectly substitutable consumption goods using their own AK technologies¹³:

$$F_j(K_t^j) = A_j K_t^j, \quad j \in \{d, g\}$$

Each sector has its own capital stock that evolves with logarithmic adjustment costs and Brownian shocks as follows:

$$dK_t^j = K_t^j[\alpha_j + \Gamma_j \log(1 + \theta_j i_t^j)]dt + \sigma_j K_t^j dW_t^j, \quad j \in \{d, g\}$$

For computational tractability, we redefine the state variables characterizing the capital stocks in the model and use $\log K = \log(K^d + K^g)$ and $Z = \frac{K^g}{K^d + K^g}$. By Ito's lemma, we can characterize the evolution of our two new state variables. Log total capital evolves as follows

$$\begin{aligned} d \log K_t &= (1 - Z_t)[\alpha_d + \Gamma_d \log(1 + \theta_d i_t^d)]dt + Z_t[\alpha_g + \Gamma_g \log(1 + \theta_g i_t^g)]dt \\ &\quad - \frac{1}{2}|\sigma_d(1 - Z_t)|^2 dt - \frac{1}{2}|\sigma_g Z_t|^2 dt + (1 - Z_t)\sigma_d dW_t^d + Z_t\sigma_g dW_t^g, \end{aligned}$$

¹³While the AK model has in part been rejected by Bernanke and Gürkaynak (2001) and others based on cross-equation restrictions relating the output-capital ratio to the sensitivity of growth to the saving rate, they also note that “the key prediction of the model that the saving rate (rate of capital accumulation) is important for explaining the growth as well as the level of per capita output seems to hold considerable validity.” This latter implication is the relevant feature for our analysis. We therefore see the AK model as a reasonable approximation for our purposes, and leave refinements of the model using alternative production technologies for future research.

while the green-capital-to-total-capital ratio evolution is given by

$$dZ_t = Z_t(1 - Z_t)\{[\alpha_g + \Gamma_g \log(1 + \theta_g i_t^g)] - [\alpha_d + \Gamma_d \log(1 + \theta_d i_t^d)]\}dt \\ + Z_t(1 - Z_t)\{(1 - Z_t)|\sigma_d|^2 - Z_t|\sigma_g|^2\}dt + Z_t(1 - Z_t)\sigma_g dW_t^g - Z_t(1 - Z_t)\sigma_d dW_t^d.$$

There are two key differences across the different types of production technologies in the economy. First, Sector d is the only sector that generates emissions, so that E_t in our temperature evolution equation is given by

$$\mathcal{E}_t = \eta A_d K_t^d,$$

where η is the emissions intensity parameter of Sector d production. Second, Sector d is initially more productive than Sector g in that $A_d > A_g$. While Sector d is initially more productive than Sector g, we also assume that there is the potential for a “green” technology shock that would augment Sector g productivity. The arrival rate of this one-time jump in Sector g productivity is an increasing function of the aggregate knowledge capital stock in the economy R . The evolution of the aggregate knowledge stock is given by

$$dR_t = -\zeta R_t dt + \psi_0 (I_t^r)^{\psi_1} (R_t)^{1-\psi_1} dt + R_t \sigma_r dW_t^r, \quad (3.1)$$

where I^r is the level of R&D investment made to increase the knowledge capital stock, ζ is the decay rate of the knowledge stock, ψ_0 and ψ_1 capture the effectiveness of R&D investment, and σ_r is the volatility associated with the knowledge capital stock.

We allow for uncertainty about the realization of the technology shock. We assume a discrete set of possible realizations for the post-technology jump Sector g productivity $A'_g(\ell)$. The technological innovation that determines the value of $A'_g(\ell)$ is realized through a stochastic jump process. There are S absorbing states, with alternative values denoted by $\ell = L - S + 1, \dots, L$, with each state ℓ corresponding to a particular value of $A'_g(\ell)$ and each value of $A'_g(\ell)$ having an equal-weighted

likelihood of occurring. For each potential technology shock realization we assume $A'_g(\ell) \geq A_d > A_g$. The likelihood of a technological innovation depends upon the jump intensity $\mathcal{J}_g^\ell(R)$, and is assumed to be increasing in the endogenous knowledge stock state variable R . The information dynamics and framework and structure of the technological change jump are similar to the damage function jump. Once the jump is realized, the true technological change outcome of $A'_g(\ell)$ is known to the social planner. Thus, with uncertainty about the probability distribution and arrival rate of potential $A'_g(\ell)$ realizations, our structure allows for a broadly conceived uncertainty analysis of a relatively novel specification of “green” technological innovation.

3.3.4. Preferences

Finally, we assume that flow utility in the model is a log function over damaged aggregate consumption, where exponential-quadratic damages multiplicatively scale consumption. The utility function is therefore given by

$$U(\tilde{C}) = U(C/N) = \delta \log(C) - \delta \log N,$$

where δ is the subjective discount rate, the two types of consumption goods are perfectly substitutable. Investment decisions are made by optimally dividing the output net of consumption across investment into the three types of capital in the economy (dirty production capital, green production capital, and knowledge capital), so that the market clearing final output constraint is given by the relationship

$$C = A_d K^d - i^d K^d + A_g K^g - i^g K^g - i^r (K^g + K^d),$$

where $i^d = I^d/K^d$ is the dirty investment-to-dirty capital ratio, $i^g = I^g/K^g$ is the green investment-to-green capital ratio, and $i^r = I^r/(K^d + K^g)$ is the R&D investment-to-total capital ratio.

3.4. Model Solution

With the model framework laid out, we can now turn to solving the Hamilton-Jacobi-Bellman equations that characterize the solution to the Social Planner’s problem in our model. The solution

to our model is a recursive Markovian equilibrium that solves the Hamilton–Jacobi–Bellman equation characterizing the planner’s social welfare function, or value function. Therefore, the equilibrium solution is determined by optimal investment choices:

$$\{i_t^{g*}, i_t^{d*}, i_t^{r*}\}$$

that maximize the planner’s discounted, lifetime expected utility as functions of the states

$$\{\log K_t, Z_t, Y_t, \log R_t, \log N_t\}.$$

These optimal controls must satisfy the evolution equations for the state variables, as well as the market clearing conditions given in the model set-up above and the first order conditions from the HJB equations we outline below. To arrive at the full model solution, we must derive solutions sequentially working from the post-technology jump, post-damage-function-jump state back to the pre-technology jump, pre-damage-function-jump state, accounting for the different possible combinations of intermediate states, which are the pre-technology jump, post-damage-function-jump state and the post-technology jump, pre-damage-function-jump state. Given the complexity and scope of the various HJB equations we must solve, we omit these details from the main text and direct the reader to Appendices B.1 – B.3 for complete model and HJB equation details.

3.5. Numerical Results

We next present and discuss the numerical model solution results, which are derived using the numerical algorithm outlined above. Before getting into the results, we briefly outline a few details for completeness regarding model assumptions, functional forms, and parameter values. After presenting and discussing the numerical results, we discuss details about how we validate our neural-network-based solutions.

3.5.1. Functional Forms, Assumptions, and Parameter Values

First, for tractability we consider the case of independent Brownian shocks, i.e.,

$$\begin{aligned} 0 &= \sigma'_d \sigma_g = \sigma'_d \varsigma = \sigma'_d \sigma_\kappa, \\ 0 &= \sigma'_g \varsigma = \sigma'_g \sigma_\kappa, \\ 0 &= \varsigma' \sigma_\kappa. \end{aligned}$$

Second, we use the following functional forms for our jump arrival rates. For the technology change jump, we assume an arrival rate that is affine in the knowledge stock:

$$\mathcal{J}_g^\ell(R) = \frac{1}{S}R$$

The knowledge stock variable is scaled such that R is in arrival rate units, and is chosen based on expected green policy implementation timelines proposed by various countries. Section B.4 provides further details about this calibration choice.

The damage jump intensity $\mathcal{J}_d^\ell(y)$ follows the arrival rate from Barnett et al. (2022):

$$\begin{aligned} \mathcal{J}_n^\ell(y) &= \left(\frac{1}{L-S} \right) \mathcal{J}_n(y), \quad \ell = 1, \dots, L-S, \\ \mathcal{J}_n(y) &= r_1 \left(\exp \left(\frac{r_2}{2} (y - \underline{y})^2 \right) - 1 \right) \mathbf{1}_{y \geq \underline{y}} \end{aligned}$$

Figure 3.3 shows that the arrival rate is an increasing function temperature anomaly y . The calibration is such that the probability of a jump occurring by $y = 2$ is essentially one.

Finally, we briefly outline the parameter values chosen for our numerical example. Further details provided in the appendix. Our economic parameters follow from Eberly and Wang (2009), whose model is a special case of our setting without climate change and technological innovation. In particular, δ is within the range of values for the subjective discount rate used in the macroeconomics




Figure 3.3: Intensity function, $r_1 = 1.5$ and $r_2 = 2.5$. With this intensity function, the probability of a jump at an anomaly of 1.6 is approximately .02 per annum, increasing to about .08 per annum at an anomaly of 1.7, increasing further to approximately .18 per annum at an anomaly of 1.8 and then to about one third per annum when the anomaly is 1.9.

and asset pricing literature. The values of (Γ_d, θ_d) and (Γ_g, θ_g) are set to be consistent with empirical measures of Tobin’s q and economic growth from BEA and World Bank databases. The values of A_d and A_g generate output consistent with World Bank World GDP values. The values of σ_d and σ_g match the time series of GDP from the World Bank database.¹⁴ We construct a set of post-technological innovation values of green productivity, $\{A'_g(\ell)\}$, using the empirical estimates of incremental and breakthrough technology steps from Acemoglu et al. (2016).

The parameters related to the knowledge stock were determined as follows. While the knowledge stock depreciation parameter ζ and R&D investment cost parameter ψ_1 are chosen in part for simplicity and computational tractability, together with the choice of the R&D investment cost parameter ψ_0 , our values are set to be consistent with estimates for the returns to R&D investment by Lucking et al. (2019) and Bloom et al. (2019), to generate R&D investment values consistent with estimates by Stine (2008) and Bloom et al. (2019) of major historical U.S. R&D investment

¹⁴We differ from Eberly and Wang (2009), who choose substantially larger values of σ_d and σ_g to produce values of the price of risk consistent with the data. Given the focus of our analysis, which includes accounting for uncertainty aversion concerns, we view our choice to match the time series standard deviation of World GDP as the appropriate empirical moment to match.

programs, and to produce an expected arrival time of a green technological innovation consistent with proposed policy timelines (between 30 and 80 years). The value of σ_r the time-series volatility of the U.S. R&D capital stock estimates in the BLS database.

The climate dynamics and climate damage parameter values are set as in Barnett et al. (2023b). Specifically, $\bar{\theta}$ is the average value across 144 climate model outcomes constructed from pulse experiments provided by Joos et al. (2013) and Geoffroy et al. (2013). The value of η is set to fit estimates of annual carbon emissions from Figueres et al. (2018) based on World Bank, EIA, and IEA estimates of the clean and dirty capital split. The value of ς comes from Barnett et al. (2022). The values of λ_1 , λ_2 , $\lambda_3(\ell)$, and \bar{y} are designed to incorporate the spread of potential climate damage outcomes based on Nordhaus (2019), Weitzman (2012), and others.

Initial values for capital, temperature anomaly, the green capital ratio, and the knowledge capital stock come from counterparts in the World Bank National Accounts data, Masson-Delmotte et al. (2021b) estimates, EIA and IEA data, and the BLS database, respectively.

3.5.2. Model Solution Results

We now discuss the computational results of our model. The results shown are simulation pathway outcomes based on the solutions to the HJB equations and are initialized at today's values of the state variables and shown out to 60 years. The results provided are for the following specifications related to model uncertainty:

- Misspecification over climate, green and dirty capital, and knowledge capital dynamics
- Misspecification over the technological change and damage curvature jumps
- Five post-jump damage model realizations: $\lambda_3(\ell) \in \{0, \frac{1}{12}, \frac{1}{6}, \frac{1}{4}, \frac{1}{3}\}$
- Three post-jump green technology realizations: $A'_g(\ell) \in \{0.120, 0.128, 0.136\}$
- Comparison across uncertainty parameters $\xi \in \{0.1, 0.3, \infty\}$

We focus on the simulated pathways from the pre-damage jump and pre-technology jump state for emissions, R&D investment, green and dirty capital investment, as well the distribution of climate, damage, and technology models, and the probabilities of a damage or technology jump occurring. We examine each of these results across different uncertainty aversion parameters to highlight the implications of increased model uncertainty on the optimal policy choices by our planner. For the simulated pathways, the red lines, represent the uncertainty neutral case when $\xi = \infty$, green lines represent the less uncertainty averse case when $\xi = 0.3$, and the blue lines represent the more uncertainty averse case when $\xi = 0.1$. For the model probabilities, the red histogram bars represent the uncertainty neutral case when $\xi = \infty$, and the blue histogram bars represent the less ($\xi = 0.3$) or more ($\xi = 0.1$) uncertainty averse as denoted in the figure.

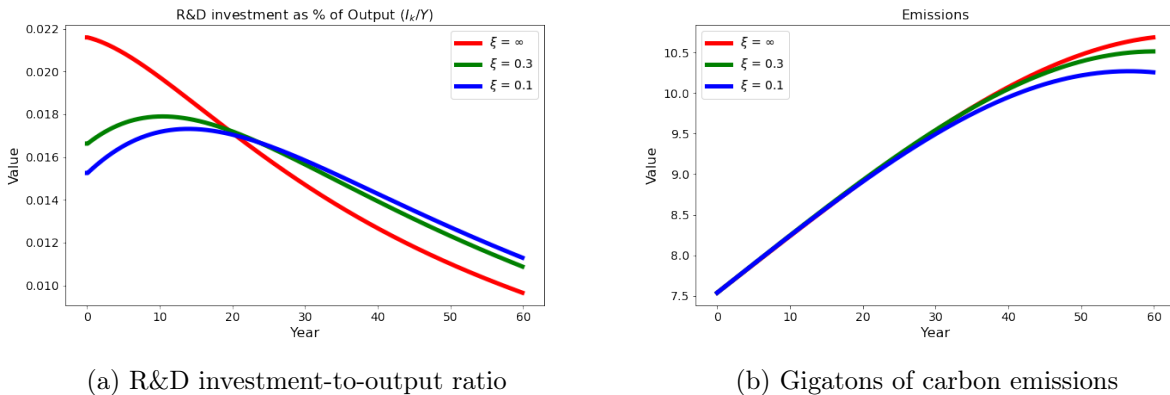


Figure 3.4: Simulated pathways of R&D investment and emissions based on the numerical solutions under different uncertainty penalty configurations. Panel (a) shows the pathway for R&D investment as a fraction of total output and Panel (b) shows the pathway for carbon emissions in gigatons of carbon (GtC). For each panel, results for three cases of model uncertainty are shown: $\xi = \infty$ (red line), $\xi = 0.3$ (green line), and $\xi = 0.1$ (blue line). The trajectories are simulated under the baseline probabilities abstracting from the intrinsic randomness.

Optimal R&D, Emissions, and Investment Choices

Figures 3.4 - 3.5 provide the main economic policy outcomes of interest. We begin by discussing Figure 3.4, which shows the pathways for optimal R&D investment and emissions over time for different cases of uncertainty aversion. In the left panel, we show R&D investment as a percent of total output. Note first that the magnitude of R&D is fairly substantial in each case, ranging between 1.5% and 2.2% of total output initially. For comparison's sake, Stine (2008) notes that

expenditure on R&D directed towards the Manhattan, Apollo, and the Federal Energy Technology programs reach magnitudes between 0.5% and 1.0%, with total R&D levels reaching over 2%. We can also see concerns about model uncertainty lead to substantial impacts to the choice of R&D investment, though perhaps in somewhat surprising ways. Initially, the more uncertainty averse is the planner the less R&D investment they make. As we will see later, this is connected in part to their decisions related to investment in the two other types of capital stock. However, over time the more uncertainty averse the planner the higher is the R&D investment. The R&D investment drops off fairly dramatically for the uncertainty neutral case, starting at 2.2% of output and dropping down to about 1.0% of output after 60 years. In the two uncertainty averse cases, though the R&D investment starts between 1.5% and 1.7% of GDP, it maintains a similar level, decreasing to about 1.2% and 1.1% after 60 years, with the highest values coming from the most uncertainty averse case starting at around 20 years into the simulated pathway.

The path pathway for emissions, shown in the right panel, are far less impacted by model uncertainty concerns. In each case of uncertainty aversion, the pathways start out at the same value and remain essentially identical out to 35 or 40 years. Also, in each case the pathways are increasing over time, though they each flatten out and will begin to decrease as they reach 60 years into the simulated pathway. An important reason for the similarity and persistence in increasing pathways of emissions in our framework is related to the fact that emissions can only be reduced by reducing the dirty capital stock. This difficulty in reducing emissions is a feature of the model structure, as we have focused on a framework aimed at addressing concerns by energy economists and others that reducing emissions is likely to be far more difficult than put forward by many existing modelling frameworks¹⁵. Another contributing factor, derived endogenously from the model solution, is that uncertainty implications appear to place greater weight on emphasizing R&D investment as compared to emissions reductions as the optimal policy. While this obviously depends upon the structure and parameterization imposed in the model, we will highlight why this is the robustly optimal policy choice for the uncertainty averse planner in our setting when we discuss the results decomposing the optimal uncertainty implications later in our analysis.

¹⁵See Reguant (2021) for one example of existing work expressing such concerns.

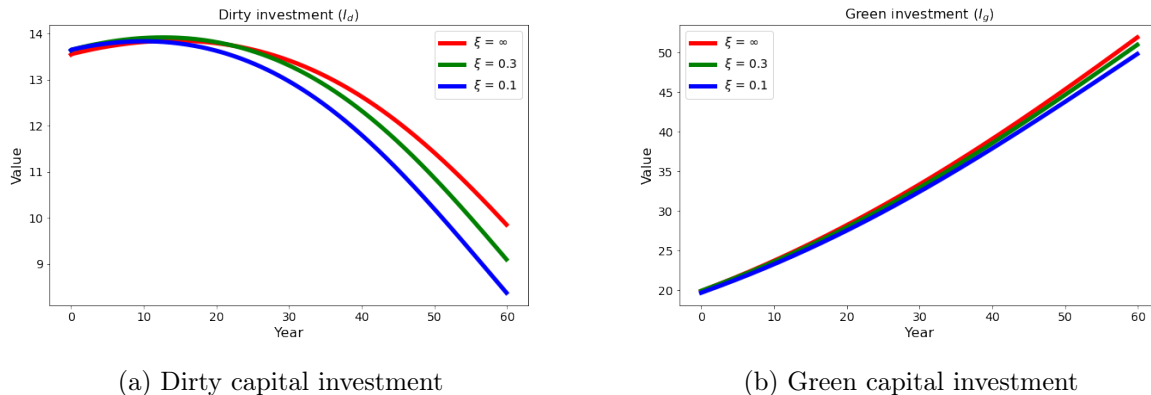


Figure 3.5: Simulated pathways of investment in dirty and green capital stocks based on the numerical solutions under different uncertainty penalty configurations. Panel (a) shows the level of investment in the dirty capital stock and Panel (b) shows the level of investment in the green capital stock. For each panel, results for three cases of model uncertainty are shown: $\xi = \infty$ (red line), $\xi = 0.3$ (green line), and $\xi = 0.1$ (blue line). The trajectories are simulated under the baseline probabilities abstracting from the intrinsic randomness.

Next, we examine the pathways for the optimal level of investment in dirty and clean capital over time for different cases of uncertainty aversion, shown in Figure 3.5. The simulated pathways for dirty capital investment are shown in the left panel, while the simulated pathways for green capital investment are shown in the right panel. Comparing outcomes across the uncertainty aversion cases and across the two types of investment, we see three key takeaways from the planner’s optimal policy choices as it relates to the two types of capital investment. First, comparing across the two types of investment, the level of investment in green capital is substantially higher than in dirty capital. At the beginning of the simulation pathway, the green investment level is about 50% higher, \$20 trillion for all cases of green investment versus roughly \$13.5 trillion for each case of dirty investment. Yet, by the end of the pathway, investment in green capital is nearly four to five times higher than for dirty capital, between \$45-\$50 trillion for green investment as compared to \$8-\$10 trillion for dirty capital. This latter comparison also highlights the second key observation related to investment across different capital stocks in that green investment persistently increases during the simulation pathway, whereas dirty investment tapers off after around 10 years and then decreases over the remaining horizon.

The final investment result comes from comparing across uncertainty averse cases, where we see that uncertainty concerns lead to limited impacts for both dirty and green capital investment. Initially, dirty investment is slightly higher for the more uncertainty averse case, whereas for green investment the values are essentially the same. However, in each case increasing uncertainty aversion leads to reduced investment in each type of capital stock over time. Dirty investment decreases more rapidly and green investment increases at a lower rate the higher is uncertainty aversion. Both of these effects are relatively modest when compared to the impact of uncertainty aversion on R&D investment. These results highlight an interesting investment trade-off being made by the planner. Initially, the higher the uncertainty aversion the lower is R&D investment and the higher is dirty capital investment. The planner invests in the higher productivity dirty capital earlier on, and then as green capital grows and climate consequences intensify, the planner shifts resources to R&D and green capital investment and away from dirty capital investment.

Probability Distortions from Uncertainty Aversion

To understand the underlying uncertainty mechanisms driving the robustly optimal policy choices we have just seen, we now examine the endogenous optimal probability distortions derived by the uncertainty averse planner in our framework. Figures 3.6 - 3.9 show the probability distortions used by the planner in making robustly optimal policy choices for the climate model, the damage model, the technology model, and the technology and damage jump arrivals.

We start by examining the probability distortions for the climate model in Figure 3.6. The left panel compares the baseline probabilities in red to the distorted probabilities for the less averse case in blue, while the right panel compares the baseline probabilities in red to the distorted probabilities for the more averse case in blue. In each case, the distortion effect is fairly modest, with a slight adjustment leading to less weight in the left tail where the least severe climate sensitivity models are and more weight in the right tail where the most severe climate sensitivity models are. The effect is only modestly more pronounced for the increased level of uncertainty aversion. So while there is some incentive for the planner to act to protect against future climate change impacts, the effect is clearly limited.

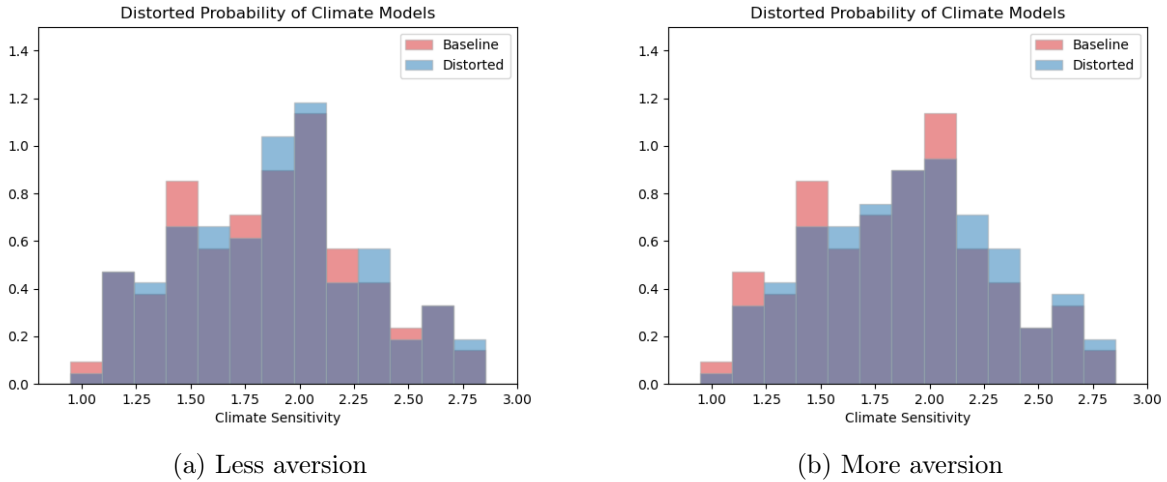


Figure 3.6: Undistorted and distorted climate model distributions. The left plot shows the undistorted and distorted distributions with less misspecification aversion ($\xi = 0.3$). The right plot shows the undistorted and distorted distributions with more misspecification aversion ($\xi = 0.3$). The histograms are calculated at year 40, with the trajectories leading up to year 40 simulated under the baseline probabilities and abstracting from the intrinsic randomness.

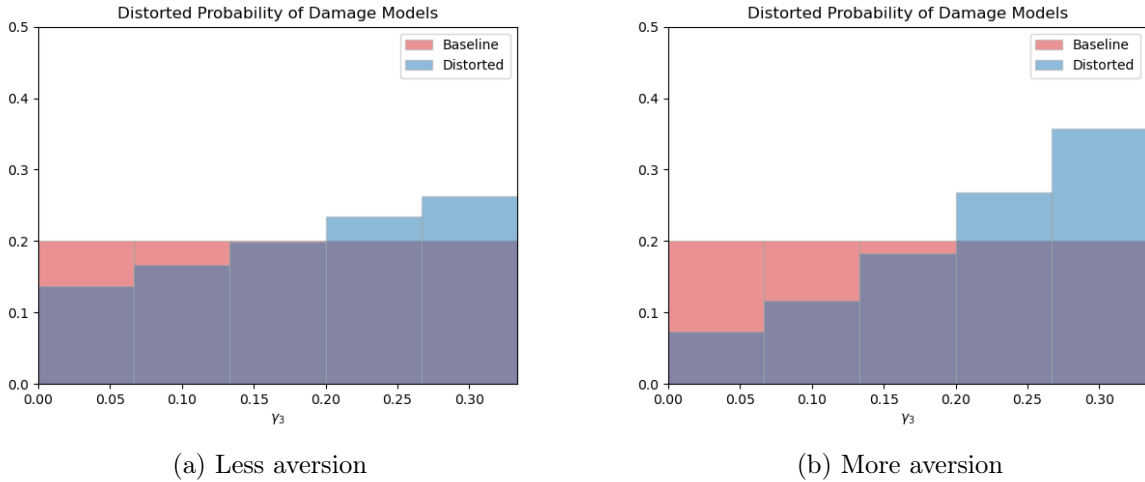


Figure 3.7: Undistorted and distorted damage model distributions. The left plot shows the undistorted and distorted distributions with less misspecification aversion ($\xi = 0.3$). The right plot shows the undistorted and distorted distributions with more misspecification aversion ($\xi = 0.3$). The histograms are calculated at year 40, with the trajectories leading up to year 40 simulated under the baseline probabilities and abstracting from the intrinsic randomness.

Figure 3.7 shows the probability distortions for the damage model. The impact of uncertainty aversion on the probabilities of the damage models use for making robust decisions by the planner

is similar to what we saw for the climate models, though effect is somewhat more pronounced along this channel. With uncertainty aversion, the planner tilts the probabilities towards the worst case damage models and away from the more favorable climate models, increasing the tilting as uncertainty aversion increases (going from the less aversion distortion in the left panel to the more aversion distortion in the right panel) moreso than we saw with the climate models. Clearly, the damage model channel adds a more substantial incentive to avoid concerns about potentially severe future climate damages than the climate model channel.

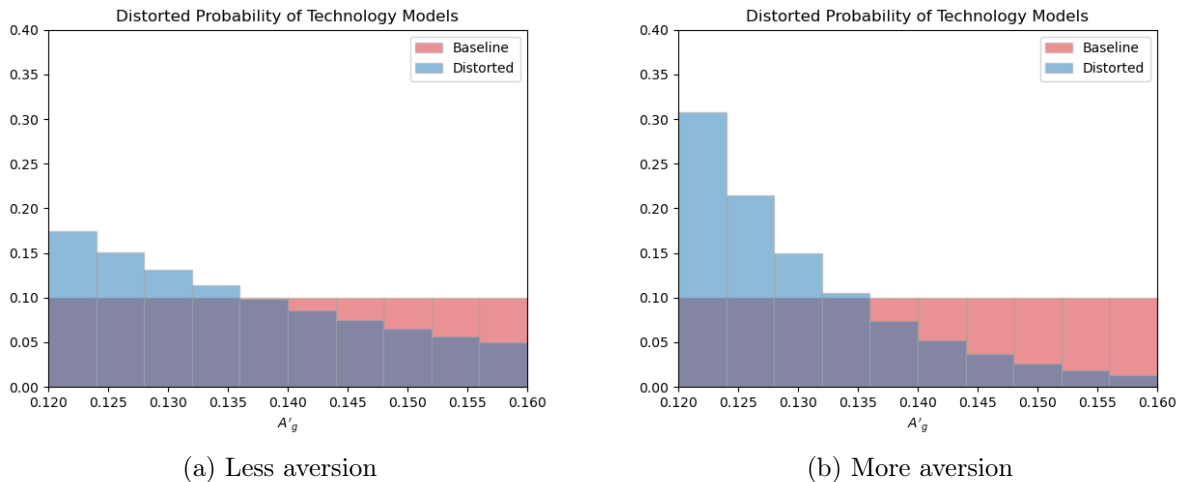


Figure 3.8: Undistorted and distorted technology model distributions. The left plot shows the undistorted and distorted distributions with less misspecification aversion ($\xi = 0.3$). The right plot shows the undistorted and distorted distributions with more misspecification aversion ($\xi = 0.3$). The histograms are calculated at year 40, with the trajectories leading up to year 40 simulated under the baseline probabilities and abstracting from the intrinsic randomness.

The impact of uncertainty aversion on the implied probability distribution for technology models is given in Figure 3.8. In this case, uncertainty aversion shifts probability weight from the right tail to the left tail, with the effect substantially amplified as we move from less aversion, in the left panel, to more aversion in the right panel. While the direction of the probability distortions differs, the intuition is the same as with the previous two channels. The planner makes optimal policy choices based on a distorted distribution that shifts probability from the best outcomes, in this case the largest green technological innovation outcomes in the right tail, to the worst outcomes, which are the smallest green technological innovation outcomes in the left tail. However, what is immediately clear

is that the magnitude of the probability distortions are the largest for the technology probabilities as compared to the climate and damage models. Thus, the planner endogenously determines that optimally robust policy should tilt most heavily towards technological innovation concerns.

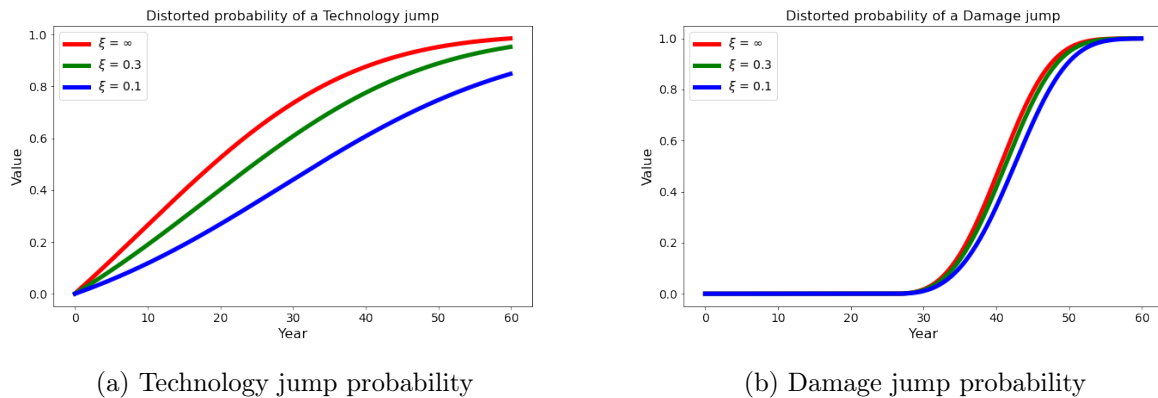


Figure 3.9: Simulated pathways of the cumulative probability of a technology or damage jump occurring based on the numerical solutions under different uncertainty penalty configurations. Panel (a) shows the distorted probability of a technology jump occurring and Panel (b) shows the distorted probability of a damage jump occurring. For each panel, results for three cases of model uncertainty are shown: $\xi = \infty$ (red line), $\xi = 0.3$ (green line), and $\xi = 0.1$ (blue line). The trajectories are simulated under the baseline probabilities abstracting from the intrinsic randomness.

The final probability distortions we examine are the pathways for the distorted jump probabilities, given in Figure 3.9. The right panel shows the distorted probability of a damage jump taking place the uncertainty neutral, less averse, and the more averse cases. It is immediately clear here that uncertainty concerns have only a very small impact to the probability of a damage jump taking place used by the planner. However, comparing across uncertainty aversion cases for the implied probability of a technology jump taking place, shown in the left panel, is a very different story. The uncertainty averse planner significantly down-weights the probability of a technology jump occurring as they incorporate more concern about model uncertainty into their robust decision making. This adjustment is one of the largest distributional impacts in relative terms, along with the distortion to the technology model probability distribution. This result highlight again that technological innovation uncertainty is the main focus of uncertainty aversion, as determine endogenously through the robust optimal decision making framework provided by dynamic decision theory. However, rather than simply reducing R&D investment due to concerns about a decreased likelihood of the

technological change taking place and that the innovation realization is treated as being increasingly more likely to be modest in magnitude, the planner initially reduces R&D investment and then over time maintains persistently higher R&D investment to try and increase the likelihood of the innovation shock occurring.

Our analysis of the impacts of uncertainty aversion, decomposed into the different probability distortion channels, highlights important effects influencing the optimal policy decisions of the planner in our model. First, the clear response of robust policy action in this setting is to avoid significant future climate change consequences. The question of how becomes clear from which probability distortions the planner finds most important in terms of welfare implications. The probability weights placed on the climate models is quite modest, with little effect in either the less or more averse case. The distortion effect is larger for the damage models, however the damage jump probability is quite limited. Most importantly from the planner’s perspective, significant probability distortions are made to the technology model weighting and the probability of a technology jump taking place. Thus, the policy response by the planner emphasizes adjustments to the R&D policy in response to uncertainty concerns, with limited adjustments to dirty and green investment as a result. The social optimality of these responses depends importantly on the perceived social payoff of such actions. Given the sticky nature of emissions given their dependence on dirty capital, the higher productivity of dirty capital relative to green capital, and the potential welfare payoff of experiencing a significant, though potentially unlikely, green technological innovation, the planner makes intersectoral and intertemporal investment trade-offs based on their uncertainty concerns which tilt toward substantial R&D investment increases in the long-run. The perceived long-run social payoff of this response outweighs alternative investment options or emissions reductions, and emphasizes the importance of R&D and technological innovation even under model uncertainty concerns.¹⁶

¹⁶In Barnett et al. (2023b), we develop a method for constructing nonlinear, stochastic impulse responses that allow us to decompose the economic mechanisms and the uncertainty contributions to these mechanisms underlying our main results into “social cash flow” and “social discount” components. We plan to apply these same methods to further develop the economic intuition for this numerical example.

3.6. Algorithm Validation

An important issue to address for our numerical solutions is the validation of the accuracy of our results. We address this in various ways, with the results shown in Figures ?? - ?. First, we examine the training loss for our neural net solutions to determine the magnitude of the HJB equation error for our solution from the DGM-PIA method. Figure B.1a shows the training error for the post-technology, post-damage jump state solution across epochs. We reach a training error of approximately 10^{-4} , suggesting our solution is reasonably accurate.¹⁷

In addition, because the post-jump model solutions require only 3 state variables, we can compare our augmented DGM-PIA algorithm results for this setting with the solution derived using more standard finite difference methods.¹⁸ Figure B.1b shows the neural net and finite difference solutions of the value function and optimal clean and dirty investment for different values of R , and values of $\log K$ and Y fixed at the midpoints of our state space range ($\log K = 5.5, Y = 1.5$). We see that the two solutions methods are consistent.

3.7. Conclusion

We develop and implement a deep learning-based algorithm to solve high-dimensional Hamilton-Jacobi-Bellman equations that account for the rich and complex nonlinearities that are of first-order importance in dynamic climate-economic models. Specifically, we focus on settings that incorporate endogenous nonlinearities that arise from model uncertainty aversion and jump processes related to climate damages and technological innovation. Model uncertainty along multiple dimensions, and jump processes characterizing economic and climate thresholds, are essential model components for a rigorous analysis of optimal climate policy for a carbon-neutral transition. Our algorithm provides one of the first approaches that we are aware of applying deep neural networks and pseudo-state variables in a way designed to handle such numerical complexities.

¹⁷Two additional tests we plan to do to further validate our methods are the following: first, examine the maximum errors across the state space and confirm errors do not concentrate in economically interesting parts of the state space; and second, solve the model with different starting seeds for the stochastic gradient descent mini-batch sampling to ensure the consistency of the model solutions accounting for the intrinsic randomness resulting from the stochastic gradient descent mini-batch sampling used in the algorithm.

¹⁸In particular, we use the false-transient, conjugate gradient-based numerical algorithm implemented in Barnett et al. (2020) and Barnett et al. (2022) for our finite difference solutions.

We analyze a particularly relevant economic example that integrates dynamic decision theory under uncertainty into a climate-economics-innovation framework with multiple capital stocks to guide uncertainty quantification of optimal R&D investment and investment in dirty and clean production capital. We examine different forms of uncertainty, including diffusion and jump process model misspecification, as well as the implications of uncertainty from climate sensitivity, climate damages, and technological innovation. Concerns about model uncertainty feed into the endogenous equilibrium policy responses in our model by “tilting” the stochastic discounting implicitly used in constructing the marginal valuations related to the externalities in our framework. This results in first-order implications for the socially optimal outcomes in our model as a result of incorporating aversion to model uncertainty in the planner’s decision problem. In particular, uncertainty aversion amplifies the incentive to invest in technological innovation as the major policy mechanism. Our algorithm allows us to derive the necessary global solutions required for computing social valuations from our continuous-time, infinite horizon social planner’s problem that optimizes lifetime expected utility subject to model uncertainty concerns with numerous endogenous and potentially non-stationary state variables.

There are various dimensions of future research that our work helps open the door to. Along the climate-economics dimension, extensions include settings with heterogeneity across agents, goods, technologies, and climate damages. There are also important considerations related to allowing explicitly for learning about economic and climate dynamics, as well as decentralized equilibria with micro-founded market frictions and policy mechanisms to explore. In addition, our numerical algorithm is well suited to address broader applications in economics and asset pricing, such as settings with rare-disasters, long-run risk, financial frictions, tipping points or thresholds, and other economically and scientifically rich frameworks. Finally, because our algorithm is designed to handle general high-dimensional problems that incorporate endogenous nonlinearities, we see significant potential for application of our methods to more general stochastic optimal control settings. We leave exploration of such alternative settings for future work.

APPENDIX A

DERIVATIONS OF CHAPTER 2

A.1. Firms' HJB equations and optimal controls

Computationally, I solve for the log of firms' value functions as a function of $\ln k$ instead of k . First, note that

$$\begin{aligned}
 \frac{d \ln Q}{d \ln k} &= \frac{1}{Q} Q_{\ln k} \\
 \frac{d \ln Q}{dg} &= \frac{1}{Q} Q_g \\
 \frac{d \ln Q}{d \hat{g}} &= \frac{1}{Q} Q_{\hat{g}} \\
 \frac{d^2 \ln Q}{d(\hat{g})^2} &= - \left[\frac{d \ln Q}{d \hat{g}} \right]^2 + \frac{1}{Q} Q_{\hat{g}\hat{g}}
 \end{aligned} \tag{A.1}$$

and Ito's lemma implies

$$d \ln k = (\iota_G g + \iota_B (1 - g) - \delta) dt$$

Dividing both sides of (2.26) and (2.12) by Q , together with (A.1), I obtain

$$\begin{aligned}
 r(\ln k, \hat{g}, A^n, \mathcal{G}) &= \max_{\iota_B, \iota_G, x_E} \frac{1}{Q} \text{CF}(\mathcal{G}) + \tilde{Q}_{\ln k} \mu_{\ln k} + \tilde{Q}_g \mu_g g + \tilde{Q}_{\hat{g}} (\hat{g} - 1) \left(\mu_{\hat{k}_B} - \mu_k + \sigma_{\hat{k}_B} x_E t \right) \\
 &+ \frac{1}{2} \left(\frac{\sigma_{\hat{k}_B}}{k} \right)^2 \left(\tilde{Q}_{\hat{g}\hat{g}} + \tilde{Q}_{\hat{g}}^2 \right) + \lambda^S \left(\frac{Q_t(A^n, \mathcal{B})}{Q_t(A^n, \mathcal{G})} - 1 \right) + \lambda^A \left(\frac{Q_t(A^{-n}, \mathcal{G})}{Q_t(A^n, \mathcal{G})} - 1 \right) \\
 r(\ln k, g, A^n, \mathcal{B}) &= \max_{\iota_B, \iota_G} \frac{1}{Q} \text{CF}(\mathcal{B}) + \tilde{Q}_{\ln k} \mu_{\ln k} + \tilde{Q}_g \mu_g g + \lambda^A \left(\frac{Q(A^{-n}, \mathcal{B})}{Q(A^n, \mathcal{B})} - 1 \right)
 \end{aligned} \tag{A.2}$$

The optimal controls are determined by the first order conditions with respect to ι_B , ι_G , and x_E .

In the good state, the FOCs are

$$\begin{aligned}
\frac{1}{Q}(-k_G - \theta_G \iota_G k_G) + \tilde{Q}_{\ln k} g + \tilde{Q}_g(1-g)g + \tilde{Q}_{\hat{g}}g(1-\hat{g}) &= 0 \quad (\iota_G) \\
\frac{1}{Q}(-k_B - \theta_B \iota_B k_B) + \tilde{Q}_{\ln k}(1-g) + \tilde{Q}_g g(-(1-g)) + \tilde{Q}_{\hat{g}}(1-g)(1-\hat{g}) &= 0 \quad (\iota_B) \\
\theta_E x_E &= \tilde{Q}_{\hat{g}}(\hat{g}-1)\sigma_{\hat{k}_B} \quad (x_E)
\end{aligned}$$

In the bad state, the FOCs are

$$\begin{aligned}
\frac{1}{Q}(-k_G - \theta_G \iota_G k_G) + \tilde{Q}_{\ln k} g + \tilde{Q}_g(1-g)g &= 0 \quad (\iota_G) \\
\frac{1}{Q}(-k_B - \theta_B \iota_B k_B) + \tilde{Q}_{\ln k}(1-g) + \tilde{Q}_g g(-(1-g)) &= 0 \quad (\iota_B)
\end{aligned}$$

A.2. Kalman-Bucy filter

I work out the derivations of the Kalman-Bucy filter. Recall the dynamics of \hat{g} is given by

$$d\hat{k}_{Bt} = \left(-\tilde{\mu}_{k_B} \hat{k}_{Bt} + 2\mu_{k_B} \tilde{\mu}_{k_B} k_{Bt} \right) dt + 2\mu_{k_B} \sigma_B k_{Bt} dW_t$$

which can be rewritten as

$$\frac{d\hat{k}_{Bt}}{\hat{k}_{Bt}} = \left(-\tilde{\mu}_{k_B} + 2\mu_{k_B} \tilde{\mu}_{k_B} \frac{1-g}{1-\hat{g}} \right) dt + 2\mu_{k_B} \sigma_B \frac{1-g}{1-\hat{g}} dW_t \quad (\text{A.3})$$

with

$$\mu_{\hat{k}_B} = -\tilde{\mu}_{k_B} + 2\mu_{k_B} \tilde{\mu}_{k_B} \frac{1-g}{1-\hat{g}} \quad \text{and} \quad \sigma_{\hat{k}_B} = 2\mu_{k_B} \sigma_B \frac{1-g}{1-\hat{g}}$$

Applying Ito's lemma, I obtain

$$d\left(\frac{\hat{k}_{Bt}}{k_t}\right) = (1-\hat{g}) \left[\left(\mu_{\hat{k}_B} - \mu_k \right) dt + \sigma_{\hat{k}_B} dW_t \right]$$

Then the posterior mean of greenness, \hat{g} , is given by

$$d\hat{g} = (\hat{g} - 1) \left[\left(\mu_{\hat{k}_B} - \mu_k \right) dt + \sigma_{\hat{k}_B} dW_t \right]$$

A.3. Numerical algorithm

Algorithm 1 Solving for general equilibrium with greenwashing

- 1: **Input:** Model parameters $\theta_G, \theta_B, \theta_E, \eta, \tau, \lambda_S, \dots$
 - 2: **Initialize:**
 - 3: Create economy with parameters $A_h, A_l, D(G), D(B), \theta_G, \theta_B, \delta, \alpha_G, \alpha_B, \alpha$
 - 4: Set up 2D model grid $(g, \ln k)$ and 3D model grid $(g, \ln k, \hat{g})$
 - 5: **Stage 1: Solve 2D model (2.19) and compute required rate of return**
 - 6: **while** not converged **and** iterations $< T$ **do**
 - 7: Compute derivatives $Q_g^H, Q_{\ln k}^H$
 - 8: Update controls ι_G, ι_B
 - 9: Update coefficients in HJB equation
 - 10: Solve linear system for value function Q^H
 - 11: Solve Kolmogorov Forward Equation
 - 12: Calculate equilibrium quantities Y, I_G, I_B, C
 - 13: Compute consumption jump j^C
 - 14: Compute required rate of return $r + \lambda^S j^C j^S$
 - 15: Check convergence
 - 16: **end while**
 - 17: **Stage 2: Initialize 3D model**
 - 18: Extend 2D solution to 3D grid
 - 19: Set initial guess for $x_E = 0$
 - 20: **Stage 3: Solve 3D model with Greenwashing**
 - 21: **while** not converged **and** iterations $< T$ **do**
 - 22: Compute derivatives $Q_g, Q_{\ln k}, Q_{\hat{g}}, Q_{\hat{g}\hat{g}}$
 - 23: Update controls ι_G, ι_B, x_E
 - 24: Update drifts $\mu_g, \mu_k, \mu_{\hat{g}}$
 - 25: Update coefficients in extended HJB equation
 - 26: Solve linear system for value function
 - 27: Check convergence
 - 28: **end while**
 - 29: **Stage 4: Simulate Transition Paths**
 - 30: Initialize N particles at steady state
 - 31: **for** $t = 1$ to T **do**
 - 32: Update state variables using optimal controls
 - 33: Record economic quantities of interest
 - 34: **end for**
 - 35: **Output:** Value functions, optimal controls, transition paths
-

APPENDIX B

DERIVATIONS OF CHAPTER 3

B.1. Example Model HJB Equations

B.1.1. Post Damage and Technology Jumps HJB Equation

The first jump-state we solve is for the post-technology jump, post-damage-function-jump state. We denote this value function by $V^{(\ell, \ell')}(K_d, K_g, \hat{Y}, \log N)$, indicating that we are at a given realization of $\lambda_3(\ell)$ and in the second technology state with green capital productivity $A'_g(\ell')$. We guess-and-verify an analytical simplification of the value function given by

$$V^{(\ell, \ell')}(K_d, K_g, \hat{Y}, \log N) = v^{(\ell, \ell')}(\log K, Z, \hat{Y}) - \log N_t.$$

The remaining HJB equation characterizing the simplified value function is given by:

$$\begin{aligned} 0 = & \max_{i^g, i^d} \min_h \delta \log([A_d - i^d](1 - Z) + [A'_g(\ell') - i^g]Z) + \delta \log K - \delta v^{(\ell, \ell')} \\ & + \left((1 - Z)\phi_d(i^d) + Z\phi_g(i^g) - \frac{\sigma_d^2(1 - Z)^2 + \sigma_g^2 Z^2}{2} \right) v_{\log K}^{(\ell, \ell')} + \frac{\sigma_d^2(1 - Z)^2 + \sigma_g^2 Z^2}{2} v_{\log K, \log K}^{(\ell, \ell')} \\ & + \left(\phi_g(i^g) - Z\sigma_g^2 - \phi_d(i^d) + (1 - Z)\sigma_d^2 \right) Z(1 - Z)v_Z^{(\ell, \ell')} + \frac{1}{2}Z^2(1 - Z)^2(\sigma_g^2 + \sigma_d^2)v_{ZZ}^{(\ell, \ell')} \\ & + \left(\left(v_{\log K}^{(\ell, \ell')} - Zv_Z^{(\ell, \ell')} \right) (1 - Z)\sigma_d + \left(v_{\log K}^{(\ell, \ell')} + (1 - Z)v_Z^{(\ell, \ell')} \right) Z\sigma_g \right) \cdot h \\ & + [-Z(1 - Z)^2\sigma_d^2 + Z^2(1 - Z)\sigma_g^2] v_{\log K, Z}^{(\ell, \ell')} \\ & + \left(\left[v_y^{(\ell, \ell')} - [\nabla D](y; \ell) \right] (\bar{\theta} + \varsigma \cdot h)\eta\mathcal{E}_t + \frac{1}{2}[\nabla^2 D](y; \ell)|\varsigma|^2(\eta\mathcal{E}_t)^2 \right) + \frac{|\varsigma|^2(\eta\mathcal{E}_t)^2}{2} v_{yy}^{(\ell, \ell')} + \xi \frac{|h|^2}{2} \end{aligned}$$

Taking first order conditions with respect to i^d and i^g , we end up with the following two expressions for optimal investment choices

$$\begin{aligned} \delta([A_d - i^d](1 - R) + [A'_g(\ell') - i^g]Z)^{-1} &= \Gamma_d \theta_d (1 + \theta_d i^d)^{-1} [v_{\log K}^{(\ell, \ell')} - Zv_Z^{(\ell, \ell')}], \\ \delta([A_d - i^d](1 - Z) + [A'_g(\ell') - i^g]Z)^{-1} &= \Gamma_g \theta_g (1 + \theta_g i^g)^{-1} [v_{\log K}^{(\ell, \ell')} + (1 - Z)v_Z^{(\ell, \ell')}]. \end{aligned}$$

Each expression equates the marginal benefit of an additional unit of a given type of capital to the marginal utility of consumption, which is the utility loss of forgoing consumption for an additional unit of investment.

B.1.2. Intermediate Jump State HJB Equations

The next two HJB equations come from the intermediate jump states. The first case is post-technology jump realization, pre-damage function jump realization. We denote this value function by $V^{(\ell')}(K_d, K_g, Y, \log N)$, indicating that we have not had a realization of $\lambda_3(\ell)$, but we are in the second technology state with green capital productivity $A'_g(\ell')$. Applying the same analytical simplification as before, the value function in this stage is given by

$$V^{(\ell')}(K_d, K_g, Y, \log N) = v^{(\ell')}(\log K, Z, Y) - \log N.$$

The simplified HJB equation in this case includes two adjustments relative to the post-technology, post-damage jump case: a term capturing the evolution of damages, and a term capturing the expectation of the damage function jump, given respectively by

$$- \left([\nabla D](y) \bar{\theta} \eta \mathcal{E}_t + \frac{1}{2} [\nabla^2 D](y) |\varsigma|^2 (\eta \mathcal{E}_t)^2 \right) + \sum_{\ell=1}^{L-S} \mathcal{J}_n^\ell(y) (v^{(\ell, \ell')} - v^{(\ell')}).$$

Expressions for the first order conditions with respect to i^d and i^g are unchanged from the post-damage, post-technology jump case, equating the marginal benefit of an additional unit of a given type of capital to the utility loss of forgoing consumption for an additional unit of investment, which is the marginal utility of consumption.

The second intermediate case is the pre-technology jump, post-damage function jump setting. We denote this value function by $V^{(\ell)}(K_d, K_g, Y, \log R, \log N)$, and applying the same analytical simplification as before, the value function in this stage is given by

$$V^{(\ell)}(K_d, K_g, \hat{Y}, \log R, \log N) = v^{(\ell)}(\log K, Z, \hat{Y}, \log R) - \log N.$$

Relative to the post-technology, post-damage jump case, we introduce terms for the following features: the evolution of the knowledge capital accounting for the possibility of R&D investment to increase the knowledge capital stock, and the expectation of the green technology jump realization altering green capital productivity to $A'_g(\ell')$:

$$\left(\psi_r(i_r) - \frac{1}{2}|\sigma_r|^2\right)v_{\log R}^{(\ell)} + \frac{|\sigma_r|^2}{2}v_{\log R, \log R}^{(\ell)} + \sum_{\ell'=L-S+1}^L \mathcal{J}_g^{\ell'}(R)[v^{(\ell, \ell')} - v^{(\ell)}].$$

In addition, allowing for R&D investment alters the output constraint in this setting, so that our flow utility term net of climate damages is now given by

$$\delta \log([A_d - i^d](1 - Z) + [A_g - i^g]Z - i^r).$$

We now take first order conditions with respect to i^d , i^g , and i^r , which give us expressions for optimal investment choices

$$\begin{aligned} \delta([A_d - i^d](1 - Z) + [A_g - i^g]Z - i^r)^{-1} &= \phi'_d(i^d)[v_{\log K}^{(\ell)} - Zv_Z^{(\ell)}], \\ \delta([A_d - i^d](1 - Z) + [A_g - i^g]Z - i^r)^{-1} &= \phi'_g(i^g)[v_{\log K}^{(\ell)} + (1 - Z)v_Z^{(\ell)}], \\ \delta([A_d - i^d](1 - Z) + [A_g - i^g]Z - i^r)^{-1} &= \psi'_r(i^r)v_{\log R}^{(\ell)}. \end{aligned}$$

As before, these expressions equate the marginal benefit of an additional unit of a given type of capital to the marginal utility of consumption, which is the utility loss of forgoing consumption for an additional unit of investment. However, given the altered output constraint expression, the marginal utility of consumption on the left-hand side of these equations now also incorporates the fact that some output is dedicated to R&D investment instead of consumption.

B.1.3. Pre Damage and Technology Jumps HJB Equation

Finally, we can write down the value function associated with the pre-technology, pre-damage jump state. We denote this value function by $V(K_d, K_g, Y, \log R, \log N)$, and again apply the same

analytical simplification to arrive at a simplified value function

$$V(K_d, K_g, Y, \log R, \log N) = v(\log K, Z, Y, \log R) - \log N.$$

The simplified HJB equation incorporates each of the changes applied to the two intermediate jump state cases: terms accounting for the evolution of the knowledge capital stock, the expectation of the technological innovation jump realization, the evolution of damages from climate change, and the expectation of the damage function jump realization, and an altered output constraint to account for R&D investment. The first order conditions with respect to i^d , i^g , and i^k match those of the pre-technology, post-damage intermediate jump state, except that the relevant value function derivatives now pertain to v . As in each case, the FOCs equate the marginal benefit of an additional unit of a given type of capital to the marginal utility of consumption, which is the utility loss of forgoing consumption for an additional unit of investment.

B.2. Model Uncertainty Details

We now detail the introduction uncertainty aversion in the form of model misspecification across the following channels of interest in our analysis:

- Climate Sensitivity to Carbon Emissions
- Dirty, Green, and Knowledge Capital Stocks and Investments
- Climate Damage Function Severity and Arrival
- “Green” Technological Change Magnitude and Arrival

Aversion to these different channels of uncertainty are introduced into our baseline model using the tools of dynamic decision theory. Importantly, the planner in our model is wary of model uncertainty of these different forms and channels, and thus adopts a preference structure to identify potential worst-case models, constrained by a penalization function, and make optimal policy decisions that are robust to these possible outcomes. We outline the different forms of model uncertainty for

our framework in what follows, focusing specifically on misspecification concerns about jump and diffusion processes.

B.2.1. Jump Misspecification Concerns

We first consider the setting with misspecification concerns about the damage function and technological change jump channels. Before the jumps occur, the relevant contributions to the HJB equation related to the technological change and damage function jumps are

$$\sum_{\ell=L-S+1}^L \mathcal{J}_g^\ell(R)(v^{(\ell)} - v) + \sum_{\ell=1}^{L-S} \mathcal{J}_n^\ell(y)(v^{(\ell)} - v).$$

Assuming a common robustness parameter ξ across each channel, these terms are replaced by the following uncertainty-adjusted contributions

$$\begin{aligned} \min_{g^\ell, g^{\ell'}} & \sum_{\ell'=L-S+1}^L \mathcal{J}_g^{\ell'}(R)g^{\ell'}(v^{(\ell')} - v) + \sum_{\ell=1}^{L-S} \mathcal{J}_n^\ell(y)g^\ell(v^{(\ell)} - v) \\ & + \xi \left[\sum_{\ell'=L-S+1}^L \mathcal{J}_g^{\ell'}(R) \left(1 - g^{\ell'} + g^{\ell'} \log g^{\ell'}\right) + \sum_{\ell=1}^{L-S} \mathcal{J}_n^\ell(y) \left(1 - g^\ell + g^\ell \log g^\ell\right) \right]. \end{aligned}$$

The adjustment introduces probability distortions g^ℓ and $g^{\ell'}$ to the jump processes for technological change and climate damages. These distortions impact expectations about the arrival rate and distribution of post-jump outcomes related to these jumps. To constrain these distortions, the relative entropy terms

$$\xi \left[\sum_{\ell'=L-S+1}^L \mathcal{J}_g^{\ell'}(R) \left(1 - g^{\ell'} + g^{\ell'} \log g^{\ell'}\right) + \sum_{\ell=1}^{L-S} \mathcal{J}_n^\ell(y) \left(1 - g^\ell + g^\ell \log g^\ell\right) \right]$$

are introduced into preferences, penalizing distortions that are “too large” based on the uncertainty aversion parameter ξ . Note that g^ℓ and $g^{\ell'}$ are endogenous objects solved for by the decision maker by minimizing discounted lifetime expected utility, subject to the relative entropy penalties. The

FOCs resulting from this minimization objective are given by

$$\begin{aligned} g^\ell &= \exp\left(-\frac{1}{\xi}(v^{(\ell)} - v)\right), \\ g^{\ell'} &= \exp\left(-\frac{1}{\xi}(v^{(\ell')} - v)\right). \end{aligned}$$

Each of these represents probability distortions that adjust the distributions associated with the jump processes related to the economic and climate channels for which there are concerns about model uncertainty. The inclusion of these probability distortions into the planner's problem leads to robustly-altered optimal policy choices that account for the distorted likelihood of outcomes based on potential worst-case outcomes.

We note that in each of the various jump realization states, there are modifications that need to be made to these components. After the technology and damage jumps have both taken place, jump uncertainty is no longer relevant for the social planner. After the technology jump has occurred, but not the damage jump, we replace A_g with $A'_g(\ell')$ and the relevant uncertainty adjustments to the HJB equation and minimization FOC pertaining to the climate model and damage function jump are given by

$$\begin{aligned} \min_{g^\ell} \sum_{\ell=1}^{L-S} \mathcal{J}_n^\ell(y) g^\ell (v^{(\ell, \ell')} - v^{(\ell')}) + \xi \sum_{\ell=1}^{L-S} \mathcal{J}_n^\ell(y) \left(1 - g^\ell + g^\ell \log g^\ell\right) \\ g^\ell = \exp\left(-\frac{1}{\xi}(v^{(\ell, \ell')} - v^{(\ell')})\right). \end{aligned}$$

Before the technology jump has occurred, but after the damage jump, the relevant uncertainty adjustments to the HJB equation and minimization FOC pertaining to the climate model and technology jump are as follows

$$\begin{aligned} \min_{g^{\ell'}} \sum_{\ell'=L-S+1}^L \mathcal{J}_g^{\ell'}(R) g^{\ell'} (v^{(\ell, \ell')} - v^{(\ell')}) + \xi \sum_{\ell'=L-S+1}^L \mathcal{J}_g^{\ell'}(R) \left(1 - g^{\ell'} + g^{\ell'} \log g^{\ell'}\right) \\ g^{\ell'} = \exp\left(-\frac{1}{\xi}(v^{(\ell, \ell')} - v^{(\ell')})\right). \end{aligned}$$

B.2.2. Diffusion Misspecification Concerns

We now consider the setting with misspecification concerns about the climate and capital stocks diffusion channels. In the pre-technology, pre-damage function justmap state the relevant contributions to the HJB equation for the diffusion components are of the form

$$\begin{aligned} & \left((1-Z)\phi_d(i^d) + Z\phi_g(i^g) - \frac{\sigma_d^2(1-Z)^2 + \sigma_g^2 Z^2}{2} \right) v_{\log K} \\ & + \left(\phi_d(i^d) - Z\sigma_g^2 - \phi_g(i^g) + (1-Z)\sigma_d^2 \right) Z(1-Z)v_Z \\ & + \left((v_y - [\nabla D](y)) \bar{\theta} \eta \mathcal{E}_t + \frac{1}{2} \lambda_2 |\varsigma|^2 (\eta \mathcal{E}_t)^2 \right) + \left(\psi_r(i^r) - \frac{1}{2} |\sigma_r|^2 \right) v_{\log R}. \end{aligned}$$

Assuming the same common robustness parameter ξ for each channel as before, these terms are replaced by the following uncertainty-adjusted contributions

$$\begin{aligned} & \min_h \left((1-Z)\phi_d(i^d) + Z\phi_g(i^g) - \frac{\sigma_d^2(1-Z)^2 + \sigma_g^2 Z^2}{2} \right) v_{\log K} \\ & + \left(\phi_g(i^g) - Z\sigma_g^2 - \phi_d(i^d) + (1-Z)\sigma_d^2 \right) Z(1-Z)v_Z \\ & + ((v_{\log K} - Zv_Z)(1-Z)\sigma_d + (v_{\log K} + (1-Z)v_Z)Z\sigma_g) \cdot h \\ & + \left((v_y - [\nabla D](y)) (\bar{\theta} + \varsigma \cdot h) \eta \mathcal{E}_t + \frac{1}{2} \lambda_2 |\varsigma|^2 (\eta \mathcal{E}_t)^2 \right) \\ & + \left(\psi_r(i^r) - \frac{1}{2} |\sigma_r|^2 + \sigma_r \cdot h \right) v_{\log R} + \xi \frac{|h|^2}{2}. \end{aligned}$$

This adjustment introduces a drift distortion h to the diffusion processes for the different types of capitals and for temperature, disguised by the Brownian shocks. This type of drift distortion is a known result from the Girsanov theorem, and captures a distributional change represented by a likelihood ratio. To constrain this distortion, again, a relative entropy term is introduced into preferences, in this setting, the quadratic expression $\xi|h|^2/2$, penalizing distortions that are “too large” based on the uncertainty aversion parameter ξ . As with g^ℓ and $g^{\ell'}$, h is an endogenous object solved by the decision maker by minimizing discounted lifetime expected utility, subject to

the relative entropy penalties. The FOC resulting from this minimization objective is given by

$$h = -\frac{1}{\xi}[\{v_{\log K} - v_Z Z\}(1 - Z), \{v_{\log K} + v_Z(1 - Z)\}Z, v_{\log R}, v_y - [\nabla D](y)] \cdot [\sigma_d, \sigma_g, \sigma_r, \varsigma \eta \mathcal{E}_t]'$$

Given the volatilities are assumed to be independent, we can separate h into four different scalar components for each dimension of the Brownian as follows:

$$\begin{aligned} h_d &= -\frac{1}{\xi}\{v_{\log K} - v_Z Z\}(1 - Z) \cdot \sigma_d \\ h_g &= -\frac{1}{\xi}\{v_{\log K} + v_Z(1 - Z)\}Z \cdot \sigma_g \\ h_r &= -\frac{1}{\xi}v_r \cdot \sigma_r \\ h_y &= -\frac{1}{\xi}(v_y - [\nabla D](y))\eta \mathcal{E}_t \cdot \varsigma \end{aligned}$$

The inclusion into the planner's problem of this probability distortion in the form of a drift distortion leads to further robustly-altered optimal policy choices that account for the distorted likelihood of outcomes based on potential worst-case outcomes.

As with the jump misspecification concerns, in each of the various jump realization states there are modifications that need to be made to these components. In the post-technology, post-damage function jump state the remaining uncertainty implications come from the climate sensitivity channel and the green and dirty capital stocks. Therefore, in addition to replacing A_g with $A'_g(\ell')$, and denoting the realized $\lambda_3(\ell)$ value, the uncertainty-adjusted contribution to the HJB and the relevant

minimization FOC are given by

$$\begin{aligned}
& \min_h \left((1-Z)\phi_d(i^d) + Z\phi_g(i^g) - \frac{\sigma_d^2(1-Z)^2 + \sigma_g^2 Z^2}{2} \right) v_{\log K}^{(\ell, \ell')} \\
& + \left(\phi_g(i^g) - Z\sigma_g^2 - \phi_d(i^d) + (1-Z)\sigma_d^2 \right) Z(1-Z)v_Z^{(\ell, \ell')} \\
& + \left(\left(v_{\log K}^{(\ell, \ell')} - Zv_Z^{(\ell, \ell')} \right) (1-Z)\sigma_d + \left(v_{\log K}^{(\ell, \ell')} + (1-Z)v_Z^{(\ell, \ell')} \right) Z\sigma_g \right) \cdot h \\
& + \left(\left[v_y^{(\ell, \ell')} - [\nabla D](y; \ell) \right] (\bar{\theta} + \varsigma \cdot h)\eta\mathcal{E}_t + \frac{1}{2}[\nabla^2 D](y; \ell)|\varsigma|^2(\eta\mathcal{E}_t)^2 \right) + \xi \frac{|h|^2}{2}, \\
& h = -\frac{1}{\xi} [\{v_{\log K}^{(\ell, \ell')} - v_Z^{(\ell, \ell')} Z\}(1-Z), \{v_{\log K}^{(\ell, \ell')} + v_Z^{(\ell, \ell')}(1-Z)\}Z, v_y^{(\ell, \ell')} - [\nabla D](y; \ell)] \\
& \cdot [\sigma_d, \sigma_g, \varsigma\eta\mathcal{E}_t]'
\end{aligned}$$

In the post-technology, pre-damage function jump state we replace A_g with $A'_g(\ell)$ and the relevant uncertainty adjustments to the HJB equation and minimization FOC pertaining to the diffusion terms are given by

$$\begin{aligned}
& \min_h \left((1-Z)\phi_d(i^d) + Z\phi_g(i^g) - \frac{\sigma_d^2(1-Z)^2 + \sigma_g^2 Z^2}{2} \right) v_{\log K}^{(\ell')} \\
& + \left(\phi_g(i^g) - Z\sigma_g^2 - \phi_d(i^d) + (1-Z)\sigma_d^2 \right) Z(1-Z)v_Z^{(\ell')} \\
& + \left(\left(v_{\log K}^{(\ell')} - Zv_Z^{(\ell')} \right) (1-Z)\sigma_d + \left(v_{\log K}^{(\ell')} + (1-Z)v_Z^{(\ell')} \right) Z\sigma_g \right) \cdot h \\
& + \left(\left[v_y^{(\ell')} - [\nabla^2 D](y) \right] (\bar{\theta} + \varsigma \cdot h)\eta\mathcal{E}_t + \frac{1}{2}\lambda_2|\varsigma|^2(\eta\mathcal{E}_t)^2 \right) + \xi \frac{|h|^2}{2}, \\
& h = -\frac{1}{\xi} [\{v_{\log K}^{(\ell')} - v_Z^{(\ell')} Z\}(1-Z), \{v_{\log K}^{(\ell')} + v_Z^{(\ell')}(1-Z)\}Z, v_y^{(\ell')} - [\nabla^2 D](y)] \cdot [\sigma_d, \sigma_g, \varsigma\eta\mathcal{E}_t]'
\end{aligned}$$

In the pre-technology, post-damage function jump state the relevant uncertainty adjustments to the

HJB equation and minimization FOC pertaining to the diffusion terms are

$$\begin{aligned}
& \min_h \left((1-Z)\phi_d(i^d) + Z\phi_g(i^g) - \frac{\sigma_d^2(1-Z)^2 + \sigma_g^2 Z^2}{2} \right) v_{\log K}^{(\ell)} \\
& + \left(\phi_g(i^g) - Z\sigma_g^2 - \phi_d(i^d) + (1-Z)\sigma_d^2 \right) Z(1-Z)v_Z^{(\ell)} \\
& + \left(\left(v_{\log K}^{(\ell)} - Zv_Z^{(\ell)} \right) (1-Z)\sigma_d + \left(v_{\log K}^{(\ell)} + (1-Z)v_Z^{(\ell)} \right) Z\sigma_g \right) \cdot h \\
& + \left(\left[v_y^{(\ell)} - [\nabla D](y; \ell) \right] (\bar{\theta} + \varsigma \cdot h)\eta\mathcal{E}_t + \frac{1}{2}[\nabla^2 D](y; \ell)|\varsigma|^2(\eta\mathcal{E}_t)^2 \right) \\
& + \left(\psi_r(i^r) - \frac{1}{2}|\sigma_r|^2 + \sigma_r \cdot h \right) v_{\log R}^{(\ell)} + \xi \frac{|h|^2}{2}, \\
& h = -\frac{1}{\xi} [\{v_{\log K}^{(\ell)} - v_Z^{(\ell)}\}Z(1-Z), \{v_{\log K}^{(\ell)} + v_Z^{(\ell)}(1-Z)\}Z, v_{\log R}^{(\ell)}, v_y^{(\ell)} - [\nabla D](y; \ell)] \\
& \cdot [\sigma_d, \sigma_g, \sigma_r, \varsigma\eta\mathcal{E}_t]'
\end{aligned}$$

B.2.3. Full Misspecification Concerns

We have so far shown the introduction of the jump and diffusion uncertainty separately in our analysis. However, in our analysis, we incorporate each of the uncertainty channels simultaneously, allowing for broadly conceived uncertainty considerations in our results. There are explicit contributions related to diffusion uncertainty and jump uncertainty in the model. However, allowing for all of the uncertainty channels simultaneously allows for important interaction effects as the jump misspecification distortions f_m and g_j , and the diffusion misspecification distortion h , will alter the value functions and their derivatives across jump states and states of nature. These more implicit effects will influence the optimized distortions arising from uncertainty aversion, and as a result, impact the social valuations and optimal policy choices of the social planner.

B.3. Full HJB Equations

B.3.1. Post Damage and Technology Jumps Setting

The HJB equation for the post damage and technology jumps setting is given by

$$\begin{aligned}
0 = & \max_{i^g, i^d} \min_h \delta \log([A_d - i^d](1 - Z) + [A'_g(\ell') - i^g]Z) + \delta \log K - \delta v^{(\ell, \ell')} \\
& + \left((1 - Z)\phi_d(i^d) + Z\phi_g(i^g) - \frac{\sigma_d^2(1 - Z)^2 + \sigma_g^2 Z^2}{2} \right) v_{\log K}^{(\ell, \ell')} + \frac{\sigma_d^2(1 - Z)^2 + \sigma_g^2 Z^2}{2} v_{\log K, \log K}^{(\ell, \ell')} \\
& + \left(\phi_g(i^g) - Z\sigma_g^2 - \phi_d(i^d) + (1 - Z)\sigma_d^2 \right) Z(1 - Z)v_Z^{(\ell, \ell')} + \frac{1}{2}Z^2(1 - Z)^2(\sigma_g^2 + \sigma_d^2)v_{ZZ}^{(\ell, \ell')} \\
& + \left(\left(v_{\log K}^{(\ell, \ell')} - Zv_Z^{(\ell, \ell')} \right) (1 - Z)\sigma_d + \left(v_{\log K}^{(\ell, \ell')} + (1 - Z)v_Z^{(\ell, \ell')} \right) Z\sigma_g \right) \cdot h \\
& + [-Z(1 - Z)^2\sigma_d^2 + Z^2(1 - Z)\sigma_g^2] v_{\log K, Z}^{(\ell, \ell')} \\
& + \left(\left[v_y^{(\ell, \ell')} - [\nabla D](y; \ell) \right] (\bar{\theta} + \varsigma \cdot h)\eta\mathcal{E}_t + \frac{1}{2}[\nabla^2 D](y; \ell)|\varsigma|^2(\eta\mathcal{E}_t)^2 \right) + \frac{|\varsigma|^2(\eta\mathcal{E}_t)^2}{2}v_{yy}^{(\ell, \ell')} + \xi \frac{|h|^2}{2}
\end{aligned}$$

B.3.2. Pre Damage and Post Technology Jumps Setting

The HJB equation for the pre damage and post technology jumps setting is given by

$$\begin{aligned}
0 = & \max_{i^g, i^d} \min_{g^\ell, h} \delta \log([A_d - i^d](1 - Z) + [A'_g(\ell') - i^g]Z - i^r) + \delta \log K - \delta v^{(\ell')} \\
& + \left((1 - Z)\phi_d(i^d) + Z\phi_g(i^g) - \frac{\sigma_d^2(1 - Z)^2 + \sigma_g^2 Z^2}{2} \right) v_{\log K}^{(\ell')} + \frac{\sigma_d^2(1 - Z)^2 + \sigma_g^2 Z^2}{2} v_{\log K, \log K}^{(\ell')} \\
& + \left(\phi_g(i^g) - Z\sigma_g^2 - \phi_d(i^d) + (1 - Z)\sigma_d^2 \right) Z(1 - Z)v_Z^{(\ell')} + \frac{1}{2}Z^2(1 - Z)^2(\sigma_g^2 + \sigma_d^2)v_{ZZ}^{(\ell')} \\
& + \left(\left(v_{\log K}^{(\ell')} - Zv_Z^{(\ell')} \right) (1 - Z)\sigma_d + \left(v_{\log K}^{(\ell')} + (1 - Z)v_Z^{(\ell')} \right) Z\sigma_g \right) \cdot h \\
& + [-Z(1 - Z)^2\sigma_d^2 + Z^2(1 - Z)\sigma_g^2] v_{\log K, Z}^{(\ell')} \\
& + \left(\left[v_y^{(\ell')} - [\nabla D](y) \right] (\bar{\theta} + \varsigma \cdot h)\eta\mathcal{E}_t + \frac{1}{2}[\nabla^2 D](y)|\varsigma|^2(\eta\mathcal{E}_t)^2 \right) + \frac{|\varsigma|^2(\eta\mathcal{E}_t)^2}{2}v_{yy}^{(\ell')} + \xi \frac{|h|^2}{2} \\
& + \sum_{\ell=1}^{L-S} \mathcal{J}_n^\ell(y)g^\ell(v^{(\ell, \ell')} - v^{(\ell')}) + \xi \sum_{\ell=1}^{L-S} \mathcal{J}_n^\ell(y) \left(1 - g^\ell + g^\ell \log g^\ell \right).
\end{aligned}$$

B.3.3. Post Damage and Pre Technology Jumps Setting

The HJB equation for the post damage and pre technology jumps setting is given by

$$\begin{aligned}
0 = & \max_{i^g, i^d, i^r} \min_{g^{\ell'}, h} \delta \log([A_d - i^d](1 - Z) + [A_g - i^g]Z - i_r) + \delta \log K - \delta v^{(\ell)} \\
& + \left((1 - Z)\phi_d(i^d) + Z\phi_g(i^g) - \frac{\sigma_d^2(1 - Z)^2 + \sigma_g^2 Z^2}{2} \right) v_{\log K}^{(\ell)} + \frac{\sigma_d^2(1 - Z)^2 + \sigma_g^2 Z^2}{2} v_{\log K, \log K}^{(\ell)} \\
& + \left(\phi_g(i^g - Z\sigma_g^2 - \phi_d(i^d) + (1 - Z)\sigma_d^2) Z(1 - Z)v_Z^{(\ell)} + \frac{1}{2}Z^2(1 - Z)^2(\sigma_g^2 + \sigma_d^2)v_{ZZ}^{(\ell)} \right. \\
& + \left. \left(v_{\log K}^{(\ell)} - Zv_Z^{(\ell)} \right) (1 - Z)\sigma_d + \left(v_{\log K}^{(\ell)} + (1 - Z)v_Z^{(\ell)} \right) Z\sigma_g \right) \cdot h \\
& + [-Z(1 - Z)^2\sigma_d^2 + Z^2(1 - Z)\sigma_g^2] v_{\log K, Z}^{(\ell)} \\
& + \left(\left[v_y^{(\ell)} - [\nabla D](y; \ell) \right] (\bar{\theta} + \varsigma \cdot h)\eta\mathcal{E}_t + \frac{1}{2}[\nabla^2 D](y; \ell)|\varsigma|^2(\eta\mathcal{E}_t)^2 \right) + \frac{|\varsigma|^2(\eta\mathcal{E}_t)^2}{2} v_{yy}^{(\ell)} + \xi \frac{|h|^2}{2} \\
& + \left(\psi_r(i^r) - \frac{1}{2}|\sigma_r|^2 + \sigma_r \cdot h \right) v_{\log R}^{(\ell)} + \frac{|\sigma_r|^2}{2} v_{\log R, \log R}^{(\ell)} + \xi \frac{|h|^2}{2} \\
& + \sum_{\ell'=L-S+1}^L \mathcal{J}_g^{\ell'}(R) g^{\ell'} (v^{(\ell, \ell')} - v^{(\ell)}) + \xi \sum_{\ell'=L-S+1}^L \mathcal{J}_g^{\ell'}(R) \left(1 - g^{\ell'} + g^{\ell'} \log g^{\ell'} \right).
\end{aligned}$$

B.3.4. Pre Damage and Technology Jumps Setting

The HJB equation for the pre damage and technology jumps setting is given by

$$\begin{aligned}
0 = & \max_{i^g, i^d, i^r} \min_{g^\ell, g^{\ell'}, h} \delta \log([A_d - i^d](1 - Z) + [A_g - i^g]Z - i^r) + \delta \log K - \delta v \\
& + \left((1 - Z)\phi_d(i^d) + Z\phi_g(i^g) - \frac{\sigma_d^2(1 - Z)^2 + \sigma_g^2 Z^2}{2} \right) v_{\log K} + \frac{\sigma_d^2(1 - Z)^2 + \sigma_g^2 Z^2}{2} v_{\log K, \log K} \\
& + \left(\phi_g(i^g) - Z\sigma_g^2 - \phi_d(i^d) + (1 - Z)\sigma_d^2 \right) Z(1 - Z)v_Z + \frac{1}{2}Z^2(1 - Z)^2(\sigma_g^2 + \sigma_d^2)v_{ZZ} \\
& + ((v_{\log K} - Zv_Z)(1 - Z)\sigma_d + (v_{\log K} + (1 - Z)v_Z)Z\sigma_g) \cdot h \\
& + [-Z(1 - Z)^2\sigma_d^2 + Z^2(1 - Z)\sigma_g^2] v_{\log K, Z} \\
& + \left([v_y - [\nabla D](y)](\bar{\theta} + \varsigma \cdot h)\eta\mathcal{E}_t + \frac{1}{2}[\nabla^2 D](y)|\varsigma|^2(\eta\mathcal{E}_t)^2 \right) + \frac{|\varsigma|^2(\eta\mathcal{E}_t)^2}{2} v_{yy} \\
& + \left(\psi_r(i^r) - \frac{1}{2}|\sigma_r|^2 + \sigma_r \cdot h \right) v_{\log R} + \frac{|\sigma_r|^2}{2} v_{\log R, \log R} + \xi \frac{|h|^2}{2} \\
& + \sum_{\ell'=L-S+1}^L \mathcal{J}_g^{\ell'}(R)g^{\ell'}(v^{(\ell')} - v) + \sum_{\ell=1}^{L-S} \mathcal{J}_n^\ell(y)g^\ell(v^{(\ell)} - v) \\
& + \xi \left[\sum_{\ell'=L-S+1}^L \mathcal{J}_g^{\ell'}(R) \left(1 - g^{\ell'} + g^{\ell'} \log g^{\ell'} \right) + \sum_{\ell=1}^{L-S} \mathcal{J}_n^\ell(y) \left(1 - g^\ell + g^\ell \log g^\ell \right) \right].
\end{aligned}$$

B.4. Parameter Values

The economic parameter values are given in Table B.1. As noted previously, the special case of our model without climate change and technological innovation is the model given in Eberly and Wang (2009). We therefore use their relevant economic parameters as our baseline solution values. These parameters are not without justification from empirical estimates and previous modeling set-ups. The choice of δ is within the range of values for the subjective discount rate used in the macroeconomics and asset pricing literature. The values of (Γ_d, θ_d) and (Γ_g, θ_g) are chosen so that the no-climate, one capital version of the model satisfies three conditions, consistent with estimated values in the macroeconomics and asset pricing literatures and empirical values from the BEA and World Bank databases, as in Barnett et al. (2020): a Tobin's q value of 2.5 ($q = \frac{A-i}{\delta}v_{\log K} = \frac{1+\theta i}{\Gamma\theta} = 2.5$); a growth rate of capital of 1.02% ($\mathbb{E}[dK/K] = \alpha + \Gamma \log(1 + \theta i) = 0.02$); and an annual risk free rate

of 1.04% ($r = 0.014$). The values of σ_d and σ_g are chosen to match the annual percent changes in the time series of GDP from the World Bank database. The values of A_d and A_g generate output consistent with World Bank World GDP values.

To determine the values of $\{A'_g(\ell)\}$ for our technological innovation framework, we build off of the empirical estimation of Acemoglu et al. (2016). In their framework, they consider both incremental technology steps, where sector j productivity A_j increases by a factor λ such that $A'_j(\ell) = A_j \times \lambda$, and breakthrough technology steps such that $A'_j(\ell) = A_j \times \lambda^{n+1}$, where n is defined by the number of steps between the leading and lagging technology in the economy $A_i/A_j = \lambda^n$. Using their estimate of the incremental step size of $\lambda = 1.063$ and of the leading and lagging technology gap of $n = 3$, we construct our range of three potential technological innovation values as follows. First, we define the “standard breakthrough” technology value for the green sector as the middle of our range, given by

$$A'_g(2) = A_g \times \lambda^{n+1} = 0.1 \times (1.063^4) = 0.128$$

We then consider the lower end of our potential realizations as being one incremental technology step below the “standard breakthrough” value. This value is given by

$$A'_g = A_g \times \lambda^n = 0.1 \times (1.063^3) = 0.12$$

which equals the productivity of the dirty sector, and so we label it as a “catching up breakthrough” realization. Finally, the upper end of our potential realizations is set to be one incremental technology step above the “standard breakthrough” value. We call this a “major breakthrough” with the value given by

$$A'_g = A_g \times \lambda^{n+2} = 0.1 \times (1.063^5) = 0.136$$

We place an equal weighted prior probability on each of the potential realizations of the technological innovation jump.

Table B.1: Economic Parameters

Parameters	values
δ	0.025
$(\alpha_d, \Gamma_d, \theta_d, \sigma_d)$	(-0.035, 0.025, 100, 0.016)
$(\alpha_g, \Gamma_g, \theta_g, \sigma_g)$	(-0.035, 0.025, 100, 0.016)
(A_d, A_g)	(0.12, 0.10)
$\{A'_g(\ell)\}_{\ell=L-S+1, \dots, L}$	$A_g \times \lambda^{n+\ell-L+S-1}, \quad \ell = L - S + 1, \dots, L$
$(\zeta, \psi_0, \psi_1, \sigma_\kappa)$	(0, 0.10583, 0.5, 0.016)
ρ	1120

We assume the depreciation of the knowledge stock to be $\zeta = 0$ for simplicity and set the R&D investment curvature parameter as $\psi_1 = 0.5$ for computational tractability. The initial value of our knowledge stock is based on estimates from the BLS of the total US R&D stock values (scaled up to World values). We set $\psi_0 = 0.1$, guided by estimates for the returns to R&D investment by Lucking et al. (2019) and Bloom et al. (2019). Given our initial knowledge stock value, these R&D investment parameters allow our model to generate R&D investment values in line with major historical U.S. R&D investment programs as estimated by Stine (2008) and Bloom et al. (2019), and generates an expected arrival time of a green technological innovation consistent with proposed policy timelines (between 30 and 80 years). The value of σ_r is chosen to match the annual percent changes in the time series of U.S. R&D capital stock from the BLS database.

Table B.2: Climate Dynamics and Damages Parameters

Parameters	values
$\bar{\theta}$	1.86 / 1000
η	0.17
ς	$1.2 \times 1.86/1000$
(λ_1, λ_2)	(0.00017675, 2×0.0022)
$\{\lambda_3(\ell)\}_{\ell=1, \dots, L-S}$	$\{\frac{1}{3} \frac{\ell-1}{L-S-1}\}_{\ell=1, \dots, L-S}$
$(r_1, r_2, \underline{y})$	(1.5, 2.5, 1.5)
\bar{y}	2

The climate dynamics and climate damage parameter values are given in Table B.2. The values of $\theta(m)$ come from pulse experiments estimates produced by Barnett et al. (2022), based on the results from Joos et al. (2013) and Geoffroy et al. (2013), and are consistent with values reported

in Masson-Delmotte et al. (2021b). The value $\bar{\theta}$ is the average value across all climate 144 models. The value of η is chosen to match the current estimate of annual carbon emissions of 10 GtC from Figueres et al. (2018), based on World Bank estimates of World GDP and EIA/IEA estimates of the clean and dirty capital split. The value of ς matches volatility used by Barnett et al. (2022). The values of λ_1 , λ_2 , $\lambda_3(\ell)$, and \bar{y} match damage function parameters used by Barnett et al. (2022), which are designed to incorporate the spread of potential climate damage outcomes based on Nordhaus (2019), Weitzman (2012), and others in the literature.

Table B.3: State Variable Initial Values and Ranges

State variables	values
K_0	739
Z_0	0.5
Y_0	1.1
R_0	11.2
State variables	range
$\log(K)$	[4, 7]
Z	[0.01, 0.99]
Y	[0, 4]
$\log(R)$	[1, 6]

For our computations, we must also specify ranges and initial values for our state variables. These values are given in Table B.3, and are set similarly to Barnett et al. (2023b). The initial value of total capital is set as $K_0 = 739.13$ so that our initial GDP matches the 2020 World GDP value of \$85 trillion estimated by the World Bank National Accounts data given our assumed initial average aggregate productivity of $\bar{\alpha} = \frac{1}{2}(A_d + A_g) = 0.11$. The initial value of global mean temperature anomaly $Y_0 = 1.1$ degrees Celsius matches the estimated current value from Masson-Delmotte et al. (2021b). The initial value of the green capital-to-total capital ratio Z_0 is based on estimates of clean and dirty capital splits from the EIA and IEA. The initial value of knowledge stock is set to $R_0 = 11.2/1120$, which matches estimates from the BLS database of the current Total US R&D capital stock value, converted and scaled up to a global value based on the US-to-world productive capital ratio. Based on this initial value, the expected arrival time of a breakthrough green technological change without additional R&D investment is the year 2100.

B.5. Neural Nets Implementations

Below we give the pseudo-code for the deep Galerkin method - policy improvement algorithms (DGM-PIA) described in Section 3.2.5, and implementation details, including network architectures and training hyperparameters. For the DGM-PIA described in Section 3.2.5, we use feedforward neural networks with 4 hidden layers of width 32 and `tanh` activation function (except for the output layer) to approximate both the value functions and the optimal controls. At the output layer, a customized hyperbolic tangent function is used for i_g, i_d , and the `sigmoid` function is used for i_κ . To train the neural nets, we run 250000 epochs with a batch size of 32. The learning rates $\text{lr}_V = 10e^{-5}$ and $\text{lr}_\alpha = 10e^{-5}$ and we use the ADAM optimizer proposed in Kingma and Ba (2014). The training is done on Google Colab with a total runtime of 2.29 hours. For the Han and E (2016) method solution, we use feedforward neural networks with 4 hidden layers of size 32 and `tanh` activation function to approximate the optimal controls. To train the neural nets, we run $N = 100,000$ epochs with a batch size of 2^7 . The infinite horizon is approximated by $[0, 1000]$, with time step size 0.1.

Algorithm 2 The DGM-PIA algorithm for solving the generic HJB (3.2.5)

Require: a maximum number of epochs N , initial neural networks for the value function $V(\mathbf{x}; \theta_0)$ and the control $\boldsymbol{\alpha}(\mathbf{x}; \varphi_0)$, learning rates for the value function and the control lr_V , lr_α

1: **for** $n = 0, 1, 2, \dots, N$ **do**

2: Generate random samples $\{\mathbf{x}_m\}_{m=1}^M$ from the domain Ω according to ν

3: Compute the value function loss in (3.2.5) using samples $\{\mathbf{x}_m\}_{m=1}^M$:

$$L_V(\theta_n; \{\mathbf{x}_m\}_{m=1}^M) = \frac{1}{M} \sum_{m=1}^M \left[-\delta V(\mathbf{x}_m; \theta_n) + L^{\boldsymbol{\alpha}(\mathbf{x}_m; \varphi_n)} V(\mathbf{x}_m; \theta_n) + f(\mathbf{x}_m, \boldsymbol{\alpha}(\mathbf{x}_m; \varphi_n)) \right]^2$$

4: Take a gradient descent step to update θ_{n+1} :

$$\theta_{n+1} = \theta_n - \text{lr}_V \nabla_{\theta} L_V(\theta_n; \{\mathbf{x}_m\}_{m=1}^M)$$

5: Compute the control loss in (3.2.5) using samples $\{\mathbf{x}_m\}_{m=1}^M$:


$$L_\alpha(\varphi_n; \{\mathbf{x}_m\}_{m=1}^M) = -\frac{1}{M} \sum_{m=1}^M \left[L^{\boldsymbol{\alpha}(\mathbf{x}_m; \varphi_n)} V(\mathbf{x}_m; \theta_n) + f(\mathbf{x}_m, \boldsymbol{\alpha}(\mathbf{x}_m; \varphi_n)) \right]^2$$

6: Take a gradient descent step to update φ_{n+1} :


$$\varphi_{n+1} = \varphi_n - \text{lr}_\alpha \nabla_{\varphi} L_\alpha(\varphi_n; \{\mathbf{x}_m\}_{m=1}^M)$$

7: **end for**

B.6. Algorithm Validation Results



(a) HJB Error for post-technology, post-damage jump state solution



(b) Finite difference and DGM-PIA solution comparison

Model solution validation. The bottom figure plots the training error over epochs. The bottom figures plot the value function and optimal controls for the FD and NN solution in the post-technology and post-damage jump state.

BIBLIOGRAPHY

- Daron Acemoglu, Philippe Aghion, Leonardo Bursztyn, and David Hemous. The environment and directed technical change. *American economic review*, 102(1):131–66, 2012.
- Daron Acemoglu, Ufuk Akcigit, Douglas Hanley, and William Kerr. Transition to clean technology. *Journal of Political Economy*, 124(1):52–104, 2016.
- Daron Acemoglu, Ali Kakhbod, and Asuman Ozdaglar. Competition in Electricity Markets with Renewable Energy Sources. *The Energy Journal*, 38:137–155, 2017. ISSN 0195-6574. URL <https://www.jstor.org/stable/26294154>. Publisher: International Association for Energy Economics.
- Viral V Acharya, Robert Engle, and Olivier Wang. Strategic commitments to decarbonize: The role of large firms, common ownership, and governments. *Work. Pap. NYU Stern Sch. Bus. New York*, 2024.
- Yves Achdou, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. *The Review of Economic Studies*, 89(1):45–86, January 2022. ISSN 0034-6527. doi: 10.1093/restud/rdab002. URL <https://doi.org/10.1093/restud/rdab002>.
- Ali Al-Aradi, Adolfo Correia, Gabriel Jardim, Danilo de Freitas Naiff, and Yuri Saporito. Extensions of the deep galerkin method. *Applied Mathematics and Computation*, 430:127287, 2022.
- M Allen, P Antwi-Agyei, F Aragon-Durand, M Babiker, P Bertoldi, M Bind, S Brown, M Buck-eridge, I Camilloni, A Cartwright, et al. Technical summary: Global warming of 1.5 °c. an ipcc special report on the impacts of global warming of 1.5 °c above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty. Technical report, IPCC, 2019.
- Fernando E. Alvarez, Francesco Lippi, and Takis Souganidis. Price Setting with Strategic Complementarities as a Mean Field Game, July 2022. URL <https://www.nber.org/papers/w30193>.
- Gene M. Amdahl. Validity of the single processor approach to achieving large scale computing capabilities. In *Proceedings of the April 18-20, 1967, Spring Joint Computer Conference, AFIPS '67 (Spring)*, page 483–485, New York, NY, USA, 1967. Association for Computing Machinery. ISBN 9781450378956. doi: 10.1145/1465482.1465560. URL <https://doi.org/10.1145/1465482.1465560>.
- Evan W Anderson, Lars Peter Hansen, and Thomas J Sargent. A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association*, 1(1):68–123, 2003.

- Svante Arrhenius. On the influence of Carbonic Acid in the Air upon Temperatue of the Ground. *Philosophical Magazine and Journal of Science Series 5*, 41:237–276, 1896.
- Marlon Azinovic, Luca Gaegauf, and Simon Scheidegger. Deep equilibrium nets. *International Economic Review*, 63(4):1471–1525, 2022.
- Achref Bachouch, Côme Huré, Nicolas Langrené, and Huyen Pham. Deep neural networks algorithms for stochastic control problems on finite horizon: numerical applications. *Methodology and Computing in Applied Probability*, pages 1–36, 2021.
- Michael Barnett. Climate Change and Uncertainty: An Asset Pricing Perspective. *Management Science*, 69(12):7562–7584, December 2023. ISSN 0025-1909. doi: 10.1287/mnsc.2022.4635. URL <https://pubsonline.informs.org/doi/abs/10.1287/mnsc.2022.4635>. Publisher: INFORMS.
- Michael Barnett, William Brock, and Lars Peter Hansen. Pricing Uncertainty Induced by Climate Change. *The Review of Financial Studies*, 33(3):1024–1066, March 2020. ISSN 0893-9454. doi: 10.1093/rfs/hhz144. URL <https://doi.org/10.1093/rfs/hhz144>.
- Michael Barnett, William Brock, and Lars Peter Hansen. Climate change uncertainty spillover in the macroeconomy. *NBER Macroeconomics Annual*, 36(1):253–320, 2022.
- Michael Barnett, William A. Brock, Lars Peter Hansen, Ruimeng Hu, and Joseph Huang. A Deep Learning Analysis of Climate Change, Innovation, and Uncertainty, October 2023a. URL <https://papers.ssrn.com/abstract=4607233>.
- Michael Barnett, William A Brock, Lars Peter Hansen, and Hong Zhong. How should climate change uncertainty impact social valuation and policy? *Working Paper*, 2023b.
- Suleyman Basak and Domenico Cuoco. An equilibrium model with restricted stock market participation. *The Review of Financial Studies*, 11(2):309–341, 1998. ISSN 08939454, 14657368. URL <http://www.jstor.org/stable/2646048>.
- William J Baumol. Productivity growth, convergence, and welfare: what the long-run data show. *The american economic review*, pages 1072–1085, 1986.
- C. Beck, W. E, and A. Jentzen. Machine learning approximation algorithms for high-dimensional fully nonlinear partial differential equations and second-order backward stochastic differential equations. *Journal of Nonlinear Science*, 29(4):1563–1619, 2019.
- Christian Beck, Sebastian Becker, Patrick Cheridito, Arnulf Jentzen, and Ariel Neufeld. Deep splitting method for parabolic pdes. *SIAM Journal on Scientific Computing*, 43(5):A3135–A3154, 2021.
- David A. Belsley, Edwin Kuh, and Roy E. Welsch. *Regression Diagnostics: Identifying Influential Data and Sources of Collinearity*. John Wiley & Sons, February 2005. ISBN 978-0-471-72514-5.

Google-Books-ID: GECBEUJVNe0C.

- Florian Berg, Julian Kolbel, and Roberto Rigobon. Aggregate Confusion: The Divergence of ESG Ratings. *SSRN Electronic Journal*, August 2019. doi: 10.2139/SSRN.3438533. URL <https://papers.ssrn.com/abstract=3438533>. Publisher: Elsevier BV.
- Ben S Bernanke and Refet S Gürkaynak. Is growth exogenous? taking mankiw, romer, and weil seriously. *NBER macroeconomics annual*, 16:11–57, 2001.
- Nicholas Bloom, John Van Reenen, and Heidi Williams. A toolkit of policies to promote innovation. *Journal of economic perspectives*, 33(3):163–84, 2019.
- C. Marshall Boffo, R. and R. Patalano. Esg investing: Environmental pillar scoring and reporting”, 2020. URL <https://www.oecd.org/finance/esg-investing-environmental-pillar-scoring-and-reporting.pdf>.
- Patrick Bolton and Marcin Kacperczyk. Do investors care about carbon risk? *Journal of Financial Economics*, 142(2):517–549, November 2021. ISSN 0304-405X. doi: 10.1016/j.jfineco.2021.05.008. URL <https://www.sciencedirect.com/science/article/pii/S0304405X21001902>.
- Stéphanie Bouckaert, Araceli Fernandez Pales, Christophe McGlade, Uwe Remme, Brent Wanner, Laszlo Varro, Davide D’Ambrosio, and Thomas Spencer. Net zero by 2050: A roadmap for the global energy sector. Technical report, International Energy Agency, 2021.
- Victor Boussange, Sebastian Becker, Arnulf Jentzen, Benno Kuckuck, and Loïc Pellissier. Deep learning approximations for non-local nonlinear pdes with neumann boundary conditions. *arXiv preprint arXiv:2205.03672*, 2022.
- W. Brock and A. Xepapadeas. Climate change policy under polar amplification. *European Economic Review*, 94:263–282, May 2017. ISSN 0014-2921. doi: 10.1016/j.euroecorev.2017.03.003. URL <https://www.sciencedirect.com/science/article/pii/S0014292117300508>.
- William A Brock. Asset prices in a production economy. In *The economics of information and uncertainty*, pages 1–46. University of Chicago Press, 1982.
- William A Brock and Leonard J Mirman. Optimal economic growth and uncertainty: the discounted case. *Journal of Economic Theory*, 4(3):479–513, 1972.
- Barry W. Brook, Erle C. Ellis, Michael P. Perring, Anson W. Mackay, and Linus Blomqvist. Does the terrestrial biosphere have planetary tipping points? *Trends in Ecology and Evolution*, 28: 396–401, 2013. ISSN 01695347. doi: 10.1016/j.tree.2013.01.016.
- Thomas Bruckner, Igor A. Bashmakov, Yacob Mulugetta, Helena Chum, Alfredo de la Vega Navarro, James Edmonds, Andre Faaij, Bundit Fungtammasan, Amit Garg, Edgar Hertwich, Damian Honnery, David Infield, Mikiko Kainuma, Smail Khennas, Son H. Kim, Hassan B. Nimir, Keywan

- Riahi, Neil Strachan, Ryan Wiser, and Xiliang Zhang. Energy systems. In Ottmar Edenhofer, Ramon Pichs-Madruga, Youba Sokona, Ellie Farahani, Susanne Kadner, Kristin Seyboth, Anna Adler, Ina Baum, Steffen Brunner, Patrick Eickemeier, Bettina Kriemann, Jukka Savolainen, Steffen Schlömer, Christoph von Stechow, Timm Zwickel, and Jan C. Minx, editors, *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 2014.
- M. K. Brunnermeier and Y. Sannikov. Chapter 18 - Macro, Money, and Finance: A Continuous-Time Approach. In John B. Taylor and Harald Uhlig, editors, *Handbook of Macroeconomics*, volume 2, pages 1497–1545. Elsevier, January 2016. doi: 10.1016/bs.hesmac.2016.06.002. URL <https://www.sciencedirect.com/science/article/pii/S1574004816300155>.
- Markus K Brunnermeier and Yuliy Sannikov. A Macroeconomic Model with a Financial Sector. *American Economic Review*, 104(2):379–421, 2014.
- Yongyang Cai and Thomas S Lontzek. The social cost of carbon with economic and climate risks. *Journal of Political Economy*, 127(6):2684–2734, 2019.
- Yongyang Cai, Kenneth L Judd, Timothy M Lenton, Thomas S Lontzek, and Daiju Narita. Environmental Tipping Points Significantly Affect the Cost-Benefit Assessment of Climate Policies. *Proceedings of the National Academy of Sciences*, 112(15):4606–4611, 2015. doi: 10.1073/pnas.1503890112. URL <https://www.pnas.org/content/112/15/4606>.
- Dario Caldara, Jesús Fernández-Villaverde, Juan F. Rubio-Ramirez, and Yao Wen. Computing dsge models with recursive preferences and stochastic volatility. Technical Report 2012-04, Finance and Economics Discussion Series, 2012.
- Stefano Carattini, Garth Heutel, and Givi Melkadze. Climate policy, financial frictions, and transition risk. *Review of Economic Dynamics*, 51:778–794, December 2023. ISSN 1094-2025. doi: 10.1016/j.red.2023.08.003. URL <https://www.sciencedirect.com/science/article/pii/S1094202523000492>.
- G. Carleo and M. Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, 2017.
- Tamma Carleton and Michael Greenstone. Updating the United States Government’s Social Cost of Carbon. SSRN Working Paper 2021-04, University of Chicago, Becker Friedman Institute for Economics, 2021.
- René Carmona and Mathieu Laurière. Convergence analysis of machine learning algorithms for the numerical solution of mean field control and games: Ii—the finite horizon case. *The Annals of Applied Probability*, 32(6):4065–4105, 2022.
- Javier Castro. Deep learning schemes for parabolic nonlocal integro-differential equations. *Partial*

- Differential Equations and Applications*, 3(6):77, 2022.
- Simone Cerreia-Vioglio, Lars Peter Hansen, Fabio Maccheroni, and Massimo Marinacci. Making Decisions under Model Misspecification, 2021.
- Darwin Choi, Zhenyu Gao, and Wenxi Jiang. Attention to Global Warming. *The Review of Financial Studies*, 33(3):1112–1145, March 2020. ISSN 0893-9454. doi: 10.1093/rfs/hhz086. URL <https://doi.org/10.1093/rfs/hhz086>.
- John H Cochrane. Production-based asset pricing and the link between stock returns and economic fluctuations. *The Journal of Finance*, 46(1):209–237, 1991.
- Lauren Cohen, Umit G. Gurun, and Quoc Nguyen. The ESG - Innovation Disconnect: Evidence from Green Patenting, July 2023. URL <https://papers.ssrn.com/abstract=3718682>.
- Riccardo Colacito, Bridget Hoffmann, and Toan Phan. Temperature and growth: A panel analysis of the united states. *Journal of Money, Credit and Banking*, 51(2-3):313–368, 2019.
- George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314, 1989. doi: 10.1007/BF02551274.
- Melissa Dell, Benjamin F Jones, and Benjamin A Olken. Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics*, 4(3):66–95, 2012.
- Klaus Desmet and Esteban Rossi-Hansberg. Climate change economics over time and space. 2023.
- Simon Dietz and Frank Venmans. Cumulative carbon emissions and economic policy: In search of general principles. *Journal of Environmental Economics and Management*, 96:108–129, 2019. ISSN 10960449. doi: 10.1016/j.jeem.2019.04.003.
- Sybren Drijfhout, Sebastian Bathiany, Claudie Beaulieu, Victor Brovkin, Martin Claussen, Chris Huntingford, Marten Scheffer, Giovanni Sgubin, and Didier Swingedouw. Catalogue of abrupt shifts in intergovernmental panel on climate change climate models. *Proceedings of the National Academy of Sciences*, 112(43):E5777–E5786, 2015.
- Victor Duarte. Gradient-based structural estimation. *SSRN*, 2023a. URL Available at SSRN: <https://ssrn.com/abstract=, 3166273>, 2018.
- Victor Duarte, Diogo Duarte, and Dejanir Silva. Machine learning for continuous-time finance. *SSRN*, 2023a. URL <https://ssrn.com/abstract=, 3012602>, 2023.
- Ran Duchin, Janet Gao, and Qiping Xu. Sustainability or Greenwashing: Evidence from the Asset Market for Industrial Pollution, April 2022. URL <https://papers.ssrn.com/abstract=4095885>.

- Darrell Duffie. *Dynamic Asset Pricing Theory: Third Edition*. Princeton University Press, January 2010. ISBN 978-1-4008-2920-0.
- Weinan E, Jiequn Han, and Arnulf Jentzen. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Commun. Math. Stat.*, 5(4):349–380, 2017. ISSN 2194-6701. doi: 10.1007/s40304-017-0117-6. URL <https://doi.org/10.1007/s40304-017-0117-6>.
- Janice Eberly and Neng Wang. Reallocating and pricing illiquid capital. In *American Economic Review, Papers and Proceedings*, 2009.
- Janice C Eberly and Neng Wang. Reallocating and Pricing Illiquid Capital: Two Productive Trees. Technical report, SSRN Electronic Journal, 2012.
- M. Eby, K. Zickfeld, A. Montenegro, D. Archer, K. J. Meissner, and A. J. Weaver. Lifetime of anthropogenic climate change: Millennial time scales of potential CO₂ and surface temperature perturbations. *Journal of Climate*, 22(10):2501–2511, 2009. ISSN 08948755. doi: 10.1175/2008JCLI2554.1.
- A. Eftekhari, D. Pasadakis, M. Bollhöfer, S. Scheidegger, and O. Schenk. Block-enhanced precision matrix estimation for large-scale datasets. *Journal of Computational Science*, 53, 2021.
- Dalya Elmalt, Deniz Igan, and Divya Kirti. Limits to Private Climate Change Mitigation, April 2021. URL <https://papers.ssrn.com/abstract=3846150>.
- Robert F Engle, Stefano Giglio, Bryan Kelly, Heebum Lee, and Johannes Stroebel. Hedging Climate Change News. *The Review of Financial Studies*, 33(3):1184–1216, March 2020. ISSN 0893-9454. doi: 10.1093/rfs/hhz072. URL <https://doi.org/10.1093/rfs/hhz072>.
- Jesus Fernandez-Villaverde, Galo Nuno, George Sorg-Langhans, and Maximilian Vogler. Solving high-dimensional dynamic programming problems using deep learning. *Unpublished working paper*, 2020.
- Jesús Fernández-Villaverde and Isaiah Hull. Dynamic programming on a quantum annealer: Solving the rbc model. Technical Report 10500, CESifo Working Paper Series, 2023.
- Jesús Fernández-Villaverde, Galo Nuño, and Jesse Perla. Taming the Curse of Dimensionality: Quantitative Economics with Deep Learning, November 2024a. URL <https://www.nber.org/papers/w33117>.
- Jesús Fernández-Villaverde, Yang Yu, and Francesco Zanetti. Technological Synergies, Heterogeneous Firms, and Idiosyncratic Volatility, 2024b. URL <https://papers.ssrn.com/abstract=4764347>.
- Christiana Figueres, Corinne Le Quéré, Anand Mahindra, Oliver Bäte, Gail Whiteman, Glen Peters,

- and Dabo Guan. Emissions Are Still Rising: Ramp Up the Cuts. *Nature*, 564:27–30, 2018.
- Christian LE Franzke, Richard Blender, Terence J O’Kane, and Valerio Lembo. Stochastic methods and complexity science in climate research and modeling. *Frontiers in Physics*, 10:931596, 2022.
- Stephie Fried, Kevin Michael Novan, and William Peterman. The Macro Effects of Climate Policy Uncertainty, March 2021. URL <https://papers.ssrn.com/abstract=3865414>.
- Aleksandra Friedl, Felix Kubler, Simon Scheidegger, and Takafumi Usui. Deep uncertainty quantification: With an application to integrated assessment models. *Working Paper*, 2023.
- Pierre Friedlingstein. Carbon cycle feedbacks and future climate change. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 373(20140421), 2015. ISSN 1364503X. doi: 10.1098/rsta.2014.0421.
- Xavier Gabaix. Chapter 4 - Behavioral inattention. In B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, editors, *Handbook of Behavioral Economics: Applications and Foundations 1*, volume 2 of *Handbook of Behavioral Economics - Foundations and Applications 2*, pages 261–343. North-Holland, January 2019. doi: 10.1016/bs.hesbe.2018.11.001. URL <https://www.sciencedirect.com/science/article/pii/S2352239918300216>.
- L. Gaedke-Merzhäuser, J. van Niekerk, O. Schenk, and H. Rue. Parallelized integrated nested laplace approximations for fast bayesian inference. *Statistics and Computing*, 33(1):25, 2023.
- X. Gao and L.-M. Duan. Efficient representation of quantum many-body states with deep neural networks. *Nature Communications*, 8(1):662, 2017.
- O Geoffroy, D Saint-Martin, D J L Olivie, A Voldoire, G Bellon, and S Tytéca. Transient Climate Response in a Two-Layer Energy-Balance Model. Part {I}: Analytical Solution and Parameter Calibration Using CMIP5 AOGCM Experiments. *Journal of Climate*, 26(6):1841–1857, 2013.
- Stefano Giglio, Bryan T. Kelly, and Johannes Stroebe. Climate Finance, December 2020. URL <https://www.nber.org/papers/w28226>.
- Nathan P Gillett. Warming proportional to cumulative carbon emissions not explained by heat and carbon sharing mixing processes. *Nature Communications*, 14(1):6466, 2023.
- Simon Glossner. ESG Risks and the Cross-Section of Stock Returns, June 2017. URL <https://papers.ssrn.com/abstract=2980917>.
- Alessandro Gnoatto, Marco Patacca, and Athena Picarelli. A deep solver for bsdes with jumps. *arXiv preprint arXiv:2211.04349*, 2022.
- Mikhail Golosov, John Hassler, Per Krusell, and Aleh Tsyvinski. Optimal Taxes on Fossil Fuel in General Equilibrium. *Econometrica*, 82(1):41–88, 2014. doi: 10.3982/ECTA10217. URL

- <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA10217>.
- Maya Gupta, Andrew Cotter, Jan Pfeifer, Konstantin Voevodski, Kevin Canini, Alexander Mangylov, Wojciech Moczydlowski, and Alexander Van Esbroeck. Monotonic calibrated interpolated look-up tables. *The Journal of Machine Learning Research*, 17(1):3790–3836, 2016.
- J. Han, A. Jentzen, and W. E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- J. Han, C. Ma, Z. Ma, and W. E. Uniformly accurate machine learning-based hydrodynamic models for kinetic equations. *Proceedings of the National Academy of Sciences*, 116(44):21983–21991, 2019.
- Jiequn Han and Weinan E. Deep learning approximation for stochastic control problems. *Deep Reinforcement Learning Workshop, arXiv preprint arXiv:1611.07422*, 2016.
- Jiequn Han and Ruimeng Hu. Recurrent neural networks for stochastic control problems with delay. *Mathematics of Control, Signals, and Systems*, 33(4):775–795, 2021.
- Jiequn Han and Jihao Long. Convergence of the deep bsde method for coupled fbsdes. *Probability, Uncertainty and Quantitative Risk*, 5:1–33, 2020.
- Lars Peter Hansen and Jianjun Miao. Aversion to Ambiguity and Model Misspecification in Dynamic Stochastic Environments. *Proceedings of the National Academy of Sciences*, 115(37):9163–9168, 2018. doi: 10.1073/pnas.1811243115. URL <http://www.pnas.org/lookup/doi/10.1073/pnas.1811243115>.
- Lars Peter Hansen and Thomas J. Sargent. Recursive Robust Estimation and Control without Commitment. *Journal of Economic Theory*, 136(1):1–27, 2007. doi: 10.1016/j.jet.2006.06.010.
- Lars Peter Hansen and Thomas J. Sargent. Robustness and ambiguity in continuous time. *Journal of Economic Theory*, 146(3):1195–1223, May 2011. ISSN 0022-0531. doi: 10.1016/j.jet.2011.01.004. URL <https://www.sciencedirect.com/science/article/pii/S0022053111000159>.
- Lars Peter Hansen and Thomas J. Sargent. Structured ambiguity and model misspecification. *Journal of Economic Theory*, 199:105165, January 2022. ISSN 0022-0531. doi: 10.1016/j.jet.2020.105165. URL <https://www.sciencedirect.com/science/article/pii/S0022053120301587>.
- Lars Peter Hansen, Joseph Huang, Paymon Khorrami, and Fabrice Tourre. Valuation dynamics in models with financial frictions: Model solution, July 2018. URL https://github.com/OpenSourceEcon/BootCamp2018/blob/master/Talks/Hansen/Huang_Presentation.pdf. Presentation slides, Open Source Macroeconomics Bootcamp.
- Lars Peter Hansen, Balint Szoke, Lloyd S. Han, and Thomas J. Sargent. Twisted Probabilities, Uncertainty, and Prices, January 2019. URL <https://papers.ssrn.com/abstract=3314356>.

- Lars Peter Hansen, Paymon Khorrami, and Fabrice Tourre. Comparative valuation dynamics in production economies: Long-run uncertainty, heterogeneity, and market frictions. *Annual Review of Financial Economics*, 16:1–38, 2024. doi: 10.1146/annurev-financial-082123-105652. First published as a Review in Advance on September 26, 2024.
- Samuel M. Hartzmark and Kelly Shue. Counterproductive Sustainable Investing: The Impact Elasticity of Brown and Green Firms, November 2022. URL <https://papers.ssrn.com/abstract=4359282>.
- Samuel M. Hartzmark and Abigail B. Sussman. Do Investors Value Sustainability? A Natural Experiment Examining Ranking and Fund Flows. *SSRN Electronic Journal*, March 2019. doi: 10.2139/SSRN.3016092. URL <https://papers.ssrn.com/abstract=3016092>. Publisher: Elsevier BV.
- Klaus Hasselmann. Stochastic climate models part i. theory. *tellus*, 28(6):473–485, 1976.
- John Hassler, Per Krusell, and Conny Olovsson. The Consequences of Uncertainty: Climate Sensitivity and Economic Sensitivity to the Climate. *Annual Review of Economics*, 10:189–205, 2018. ISSN 19411391. doi: 10.1146/annurev-economics-080217-053229.
- Harrison Hong, Neng Wang, and Jinqiang Yang. Mitigating Disaster Risks in the Age of Climate Change. *Econometrica*, 91(5):1763–1802, 2023a. ISSN 1468-0262. doi: 10.3982/ECTA20442. URL <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA20442>. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA20442>.
- Harrison Hong, Neng Wang, and Jinqiang Yang. Welfare Consequences of Sustainable Finance. *The Review of Financial Studies*, 36(12):4864–4918, December 2023b. ISSN 0893-9454. doi: 10.1093/rfs/hhad048. URL <https://doi.org/10.1093/rfs/hhad048>.
- Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2):251–257, 1991. doi: 10.1016/0893-6080(91)90009-T.
- Solomon Hsiang, Robert Kopp, Amir Jina, James Rising, Michael Delgado, Shashank Mohan, DJ Rasmussen, Robert Muir-Wood, Paul Wilson, Michael Oppenheimer, et al. Estimating Economic Damage from Climate Change in the United States. *Science*, 356(6345):1362–1369, 2017.
- Joseph Huang, Haomin Qin, and Chun Hei Hung. Computational appendix (appendix c). In Lars Peter Hansen, Paymon Khorrami, and Fabrice Tourre, editors, *Comparative Valuation Dynamics in Production Economies: Long-Run Uncertainty, Heterogeneity, and Market Frictions*, volume 16, pages 1–38, Appendix C. 2024. doi: 10.1146/annurev-financial-082123-105652. This appendix was authored by Joseph Huang, Haomin Qin, and Chun Hei Hung.
- Côme Huré, Huyên Pham, and Xavier Warin. Deep backward schemes for high-dimensional nonlinear PDEs. *Mathematics of Computation*, 89(324):1547–1579, 2020.

- Martin Hutzenthaler, Arnulf Jentzen, Thomas Kruse, et al. On multilevel picard numerical approximations for high-dimensional nonlinear parabolic partial differential equations and high-dimensional nonlinear backward stochastic differential equations. *Journal of Scientific Computing*, 79(3):1534–1571, 2019.
- IMF. Fiscal monitor: How to mitigate climate change. international monetary fund, 2019.
- Niko Jaakkola and Frederick Van der Ploeg. Non-cooperative and cooperative climate policies with anticipated breakthrough technology. *Journal of Environmental Economics and Management*, 97: 42–66, 2019.
- JD Jenkins, EN Mayfield, ED Larson, SW Pacala, and C Greig. Mission net-zero america: The nation-building path to a prosperous, net-zero emissions economy., 2021.
- Urban J Jermann. Asset pricing in production economies. *Journal of monetary Economics*, 41(2): 257–275, 1998.
- Shaolin Ji, Shige Peng, Ying Peng, and Xichuan Zhang. Three algorithms for solving high-dimensional fully coupled fbsdes through deep learning. *IEEE Intelligent Systems*, 35(3):71–84, 2020.
- F. Joos, R. Roth, J. S. Fuglestedt, G. P. Peters, I. G. Enting, W. Von Bloh, V. Brovkin, E. J. Burke, M. Eby, N. R. Edwards, T. Friedrich, T. L. Frölicher, P. R. Halloran, P. B. Holden, C. Jones, T. Kleinen, F. T. Mackenzie, K. Matsumoto, M. Meinshausen, G. K. Plattner, A. Reisinger, J. Segschneider, G. Shaffer, M. Steinacher, K. Strassmann, K. Tanaka, A. Timmermann, and A. J. Weaver. Carbon Dioxide and Climate Impulse Response Functions for the Computation of Greenhouse Gas Metrics: A Multi-Model Analysis. *Atmospheric Chemistry and Physics*, 13(5): 2793–2825, 2013. doi: 10.5194/acp-13-2793-2013.
- Kenneth L. Judd. Approximation, perturbation, and projection methods in economic analysis. In H. M. Amman, D. A. Kendrick, and J. Rust, editors, *Handbook of Computational Economics*, volume 1, chapter 12, pages 509–585. Elsevier, 1 edition, 1996.
- Greg Kaplan, Benjamin Moll, and Giovanni L. Violante. Monetary policy according to hank. *American Economic Review*, 108(3):697–743, March 2018. doi: 10.1257/aer.20160042. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20160042>.
- Christos Karydas and Anastasios Xepapadeas. Pricing Climate Change Risks: CAPM with Rare Disasters and Stochastic Probabilities, January 2019. URL <https://papers.ssrn.com/abstract=3324499>.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- P Klibanoff, M Marinacci, and S Mukerji. Recursive {Smooth} {Ambiguity} {Preferences}. *Journal*

- of Economic Theory*, 144:930–976, 2009.
- Leonid Kogan and Dimitris Papanikolaou. Growth opportunities, technology shocks, and asset prices. *The journal of finance*, 69(2):675–718, 2014.
- Nikola B Kovachki, Samuel Lanthaler, and Andrew M Stuart. Operator learning: Algorithms and analysis. *arXiv preprint arXiv:2402.15715*, 2024.
- Howard Kung and Lukas Schmid. Innovation, growth, and asset prices. *The Journal of Finance*, 70(3):1001–1037, 2015.
- Andrea Lanteri and A. Rampini. Financing the Adoption of Clean Technology. 2023. URL <https://www.semanticscholar.org/paper/Financing-the-Adoption-of-Clean-Technology%E2%88%97-Lanteri-Rampini/7be8df938826f9d8af3f843d6a0a33da6490c261>.
- Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, March 2007. ISSN 1861-3624. doi: 10.1007/s11537-007-0657-8. URL <https://doi.org/10.1007/s11537-007-0657-8>.
- Derek Lemoine and Christian Traeger. Watch your step: optimal policy in a tipping climate. *American Economic Journal: Economic Policy*, 6(1):137–66, 2014.
- Timothy M Lenton, Hermann Held, Elmar Kriegler, Jim W Hall, Wolfgang Lucht, Stefan Rahmstorf, and Hans Joachim Schellnhuber. Tipping Elements in the Earth’s Climate System. *Proceedings of the National Academy of Sciences*, 105(6):1786–1793, 2008. doi: 10.1073/pnas.0705414105.
- David Levitan. Quick-change planet: Do global climate tipping points exist? *Scientific American*, 3 2013.
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.
- Robert S. Liptser and Albert N. Shiryaev. The Structure of Local Martingales, Absolute Continuity of Measures for Point Processes, and Filtering. In Robert S. Liptser and Albert N. Shiryaev, editors, *Statistics of Random Processes: II. Applications*, Stochastic Modelling and Applied Probability, pages 309–353. Springer, Berlin, Heidelberg, 2001. ISBN 978-3-662-10028-8. doi: 10.1007/978-3-662-10028-8_9. URL https://doi.org/10.1007/978-3-662-10028-8_9.
- Liwei Lu, Ruimeng Hu, Xu Yang, and Yi Zhu. Multi-agent relative investment games in a jump diffusion market with deep reinforcement learning algorithm. *arXiv preprint arXiv:2404.11967*, 2024.
- Robert E Lucas Jr. On the mechanics of economic development. *Journal of monetary economics*, 22(1):3–42, 1988.

- Brian Lucking, Nicholas Bloom, and John Van Reenen. Have r&d spillovers declined in the 21st century? *Fiscal Studies*, 40(4):561–590, 2019.
- Fabio Maccheroni, Massimo Marinacci, and Aldo Rustichini. Dynamic Variational Preferences. *Journal of Economic Theory*, 128(1):4–44, 2006. doi: 10.1016/j.jet.2005.12.011.
- Andrew H. MacDougall, Neil C. Swart, and Reto Knutti. The Uncertainty in the Transient Climate Response to Cumulative CO₂ Emissions Arising from the Uncertainty in Physical Climate Parameters. *Journal of Climate*, 30(2):813–827, 2017. doi: 10.1175/JCLI-D-16-0205.1. URL <http://dx.doi.org/10.1175/JCLI-D-16-0205.1>.
- Lilia Maliar, Serguei Maliar, and Pablo Winant. Deep learning for solving dynamic economic models. *Journal of Monetary Economics*, 122:76–101, 2021.
- V. Masson-Delmotte, P. Zhai, A. Pirani, S.L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews, T.K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou. Summary for policymakers - climate change 2021: The physical science basis. contribution of working group i to the sixth assessment report of the intergovernmental panel on climate change. Book section, IPCC, Cambridge, United Kingdom and New York, NY, USA, 2021a.
- V. Masson-Delmotte, P. Zhai, A. Pirani, S.L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews, T.K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B.u Zhou. Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Technical report, IPCC, 2021b.
- H Damon Matthews, Nathan P Gillett, Peter A Stott, and Kirsten Zickfeld. The proportionality of global warming to cumulative carbon emissions. *Nature*, 459(7248):829–832, 2009. ISSN 00280836. doi: 10.1038/nature08047.
- Peter Maxted. A Macro-Finance Model with Sentiment. *The Review of Economic Studies*, page rdad023, March 2023. ISSN 0034-6527. doi: 10.1093/restud/rdad023. URL <https://doi.org/10.1093/restud/rdad023>.
- McKinsey Global Initiative. The net-zero transition: What it would cost, what it could bring. mckensey global institute, 2022.
- William Nordhaus. Projections and uncertainties about climate change in an era of minimal climate policies. *American Economic Journal: Economic Policy*, 10(3):333–360, 2018. doi: 10.1257/pol.20170046.
- William Nordhaus. Economics of the disintegration of the greenland ice sheet. *Proceedings of the National Academy of Sciences*, 116(25):12261–12269, 2019.

- Andriy Norets. Estimation of dynamic discrete choice models using artificial neural network approximations. *Econometric Reviews*, 31(1):84–106, 2012.
- OECD. *Regions in industrial transition: Policies for people and places*. ORGANIZATION FOR ECONOMIC, 2019.
- Roman Olson, Ryan Sriver, Marlos Goes, Nathan M. Urban, H. Damon Matthews, Murali Haran, and Klaus Keller. A Climate Sensitivity Estimate using Bayesian Fusion of Instrumental Observations and an Earth System Model. *Journal of Geophysical Research Atmospheres*, 117(D04103): 1–11, 2012. doi: 10.1029/2011JD016620.
- Tim Palmer and Bjorn Stevens. The scientific challenge of understanding and estimating climate change. *Proceedings of the National Academy of Sciences*, 116(49):24390–24395, 2019.
- Dimitris Papanikolaou. Investment shocks and asset prices. *Journal of Political Economy*, 119(4): 639–685, 2011.
- D. Pasadakis, M. Bollhöfer, and O. Schenk. Sparse quadratic approximation for graph learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9):11256–11269, 2023.
- Lubos Pastor, Robert F. Stambaugh, and Lucian A. Taylor. Dissecting Green Returns. *SSRN Electronic Journal*, February 2022. doi: 10.2139/SSRN.3864502. URL <https://papers.ssrn.com/abstract=3864502>. Publisher: Elsevier BV.
- Lübosš Pástor and Pietro Veronesi. Uncertainty about government policy and stock prices. *The Journal of Finance*, 67(4):1219–1264, 2012. ISSN 1540-6261. doi: 10.1111/j.1540-6261.2012.01746.x. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.2012.01746.x>.
- Jonathan Payne, Adam Rebei, and Yucheng Yang. Deep learning for search and matching models. *Available at SSRN*, 2024.
- Huyen Pham, Xavier Warin, and Maximilien Germain. Neural networks-based backward scheme for fully nonlinear pdes. *SN Partial Differential Equations and Applications*, 2(1):16, 2021.
- Ramond T. Pierrehumbert. Short-Lived Climate Pollution. *Annual Review of Earth and Planetary Science*, 42:341–379, 2014.
- H. O. Pörtner, D. C. Roberts, M. Tignor, E. S. Poloczanska, K. Mintenbeck, A. Alegría, M. Craig, S. Langsdorf, S. Löschke, V. Möller, A. Okem, and B. Rama. Summary for policymakers - climate change 2022: Impacts, adaptation, and vulnerability. contribution of working group ii to the sixth assessment report of the intergovernmental panel on climate change. Book section, IPCC, Cambridge, UK, 2022. In Press.
- Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial

- differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- Mar Reguant. Comment on “climate change uncertainty spillover in the macroeconomy”. Technical report, National Bureau of Economic Research, 2021.
- Katharine L. Ricke and Ken Caldeira. Maximum Warming Occurs about One Decade After a Carbon Dioxide Emission. *Environmental Research Letters*, 9(12):1–8, 2014. ISSN 17489326. doi: 10.1088/1748-9326/9/12/124002.
- Paul DL Ritchie, Joseph J Clarke, Peter M Cox, and Chris Huntingford. Overshooting tipping point thresholds in a changing climate. *Nature*, 592(7855):517–523, 2021.
- Paul M Romer. Endogenous technological change. *Journal of political Economy*, 98(5, Part 2): S71–S102, 1990.
- K. Geert Rouwenhorst. Asset Pricing Implications of Equilibrium Business Cycle Models. In Thomas F. Cooley, editor, *Frontiers of Business Cycle Research*, pages 294–330. Princeton University Press, 1995. ISBN 978-0-691-04323-4. doi: 10.2307/j.ctv14163jx.16. URL <http://www.jstor.org/stable/j.ctv14163jx.16>.
- Ivan Rudik. Optimal climate policy when damages are unknown. *American Economic Journal: Economic Policy*, 12(2):340–73, May 2020. doi: 10.1257/pol.20160541. URL <https://www.aeaweb.org/articles?id=10.1257/pol.20160541>.
- Davor Runje and Sharath M Shankaranarayana. Constrained monotonic neural networks. *arXiv preprint arXiv:2205.11775*, 2022.
- John Rust. Numerical dynamic programming in economics. In H. M. Amman, D. A. Kendrick, and J. Rust, editors, *Handbook of Computational Economics*, volume 1, chapter 14, pages 619–729. Elsevier, 1 edition, 1996.
- Yuri F Saporito and Zhaoyu Zhang. Path-dependent deep galerkin method: A neural network approach to solve path-dependent partial differential equations. *SIAM Journal on Financial Mathematics*, 12(3):912–940, 2021.
- Parinitha Sastry, Emil Verner, and David Marques-Ibanez. Business as Usual: Bank Net Zero Commitments, Lending, and Engagement, March 2024. URL <https://papers.ssrn.com/abstract=4663933>.
- Maxime Sauzet. Projection methods via neural networks for continuous-time models. *Available at SSRN 3981838*, 2021.
- Jochen Schmittmann and Yun Gao. Green Bond Pricing and Greenwashing Under Asymmetric Information, January 2023. URL <https://papers.ssrn.com/abstract=4322389>.

- Joseph Sill. Monotonic networks. *Advances in neural information processing systems*, 10, 1997.
- J. Sirignano and K. Spiliopoulos. DGM: A deep learning algorithm for solving partial differential equations. *Journal of computational physics*, 375:1339–1364, 2018.
- Hwijae Son, Jin Woo Jang, Woo Jin Han, and Hyung Ju Hwang. Sobolev Training for Physics Informed Neural Networks. January 2021. doi: 10.48550/arxiv.2101.08932. URL <https://arxiv.org/abs/2101.08932v2>. arXiv: 2101.08932.
- Deborah D Stine. *The Manhattan Project, the Apollo program, and federal energy technology R & D programs: A comparative analysis*. Library of Congress, 2008.
- Neng Wang and Janice C. Eberly. Reallocating and Pricing Illiquid Capital: Two Productive Trees. *SSRN Electronic Journal*, November 2010. doi: 10.2139/SSRN.1915332. URL <https://papers.ssrn.com/abstract=1915332>. Publisher: Elsevier BV.
- Antoine Wehenkel and Gilles Louppe. Unconstrained monotonic neural networks. *Advances in neural information processing systems*, 32, 2019.
- David A. Weisbach and Gilbert E. Metcalf. The design of a carbon tax. *Harvard Environmental Law Review*, 33:499, 2009.
- Martin L. Weitzman. GHG Targets as Insurance Against Catastrophic Climate Damages. *Journal of Public Economic Theory*, 14(2):221–244, 2012. doi: 10.1111/j.1467-9779.2011.01539.x.
- White House. The long-term strategy of the united states: pathways to net-zero greenhouse gas emissions by 2050, 2021.
- L. Zhang, J. Han, H. Wang, R. Car, and W. E. Deep potential molecular dynamics: a scalable model with the accuracy of quantum mechanics. *Physical Review Letters*, 120(14):143001, 2018.
- Bernt Øksendal. The Filtering Problem. In Bernt Øksendal, editor, *Stochastic Differential Equations: An Introduction with Applications*, pages 79–106. Springer, Berlin, Heidelberg, 1998. ISBN 978-3-662-03620-4. doi: 10.1007/978-3-662-03620-4_6. URL https://doi.org/10.1007/978-3-662-03620-4_6.