Green portfolios*

Thomas Fischer† Frederik Lundtofte‡ Juan Carlos Parra-Alvarez§

January 2026

Abstract

We develop a long-horizon framework combining Merton's intertemporal portfolio theory with the damage-based logic of Nordhaus-type climate models. The Social Cost of Carbon (SCC) emerges as a key driver of optimal portfolio choice and equilibrium asset pricing. We derive a four-fund separation result. We find that stocks with sufficiently adverse climate effects may exhibit positive alphas relative to the CAPM. Using US stock market data from 2007–2019, we estimate the SCC and show that climate externalities lead to measurable, though moderate, adjustments in the cost of capital. Our model offers a tractable integration of environmental damages into portfolio theory.

Keywords: green investments, climate change, portfolio selection, asset pricing
JEL classification: G11, G12, Q51, Q54

*We thank Hossein Asgharian, Tony Berrada, Felix Csajka, Lieske Coumans, David Feldman, Pascal Gisclon, Christoph Hambel, Po-Hsuan Hsu, Chanik Jo, Paul Klumpes, Felix Kunz, Christoph Meinerding, Lakshmi Naaraayanan, Nhut Nguyen Hoang, Luca Taschini, Raman Uppal, Anders Vilhelmsson, the participants of the 2025 Asian Finance Association Annual Conference, the 2025 Econometric Society Interdisciplinary Frontiers (ESIF) and CEPR conference on Economics+Climate Science, the 2023 China International Conference in Finance, the 2023 FMA Europe Conference, the KWC/SNEE Conference on Sustainable Finance, the 2021 Aalborg University Business School Conference, and seminar participants at Lund University for helpful comments. Special thanks to Hossein Asgharian, Michal Dzielinski, and Zahra Hashemzadeh for invaluable help with the data. Thomas Fischer and Frederik Lundtofte gratefully acknowledge financial support from the Jan Wallander and Tom Hedelius Foundation. Thomas Fischer gratefully acknowledges financial support from FORMAS. Frederik Lundtofte gratefully acknowledges financial support from the Otto Mønsted Foundation. Juan Carlos Parra-Alvarez gratefully acknowledges financial support from CoRE. Computations in this project were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS) at the National Supercomputer Center (NSC) partially funded by the Swedish Research Council through grant agreement no. 2022-06725 (project number: NAISS 2024/5-457).

†Department of Economics, Lund University; Knut Wicksell Centre for Financial Studies; E-mail address: thomas.fischer@nek.lu.se

‡*Corresponding author.* Aalborg University Business School, Aalborg University; Danish Finance Institute; E-mail address: fsl@business.aau.dk

§Department of Economics and Business Economics, Aarhus University; CoRE, Center for Research in Energy: Economics and Markets; Danish Finance Institute; E-mail address: jparra@econ.au.dk

1. Introduction

In October 2021, ABP, one of the world's largest pension funds, announced that it would divest approximately €15 billion of fossil fuel holdings by early 2023.[1] This decision reflects a broader shift among institutional investors toward incorporating climate considerations into portfolio decisions. According to the Global Sustainable Investment Alliance, assets managed under sustainable investment strategies represented more than 30% of total professionally managed assets in the United States in 2020.[2] Yet despite the rapid growth of sustainable investing, there remains limited theoretical guidance on how investors should optimally incorporate climate externalities into portfolio decisions. Through the price mechanism, large-scale portfolio reallocations can affect firms' financing costs and thereby influence the trajectory of future emissions. Assessing a portfolio's climate impact, however, remains challenging: the mapping from emissions to economic damage is uncertain, firms are heterogeneous in their exposure and climate intensity, and equity holdings themselves shape future CO_2 emissions by financing production. These developments have sharpened the need for a rigorous framework linking portfolio choices, asset prices, and climate dynamics. This paper develops such a framework and addresses three central questions: How should one invest optimally while accounting for climate concerns? How does a firm's climate impact affect its equilibrium return when investors price carbon externalities? How does uncertainty about the climate, in turn, influence optimal portfolios and equilibrium prices?

To address these questions, we develop a unified, infinite-horizon framework that merges the dynamic portfolio logic of Merton's (1971; 1973) Intertemporal Capital Asset Pricing Model (ICAPM) with the damage-based structure of Nordhaus (1993)'s Integrated Assessment Model (IAMs). We refer to this framework as the *Merton–Nordhaus model*. The model extends the ICAPM to a modern and flexible setting with stochastic differential utility (Duffie and Epstein, 1992a,b), allowing for uncertainty regarding the impact of CO_2 emissions. In this setup, a representative investor chooses an optimal portfolio of risky assets while internalizing the effect of firms' CO_2 emissions on future production possibilities. Emissions are endogenous and affect non-financial income through a damage function, as in IAMs. Although climate concerns are not directly embedded in the investor's preferences, she rationally accounts for the effect of emissions on future income and wealth.

Within this setting, we derive expressions for optimal consumption and portfolio allocations and formally identify the *Social Cost of Carbon* (SCC) from the model's equilibrium conditions. We then show how the SCC influences investment and consumption

[1] See https://www.reuters.com/business/sustainable-business/ (retrieved August 28, 2025).

[2] See http://www.gsi-alliance.org/wp-content/uploads/2021/08/GSIR-20201.pdf While not all sustainable investments are motivated by climate concerns, these figures, together with numerous high-profile divestment announcements, indicate a growing focus on climate among investors.

choices, providing a clear link between environmental externalities and portfolio behavior. To our knowledge, this is the first paper to deliver such analytical results in a general multi-asset environment with endogenous emissions and stochastic differential utility. Noticeably, we obtain a four-fund separation result. That is, the optimal portfolio comprises four funds: i) the risk-free asset, ii) a myopic fund (Merton, 1969, 1971), iii) a fund reflecting climate externalities, and iv) a fund hedging the impact of CO_2 emissions on future production. This decomposition provides a transparent link between climate risk, externalities, and portfolio choice.

After deriving the investor's optimal portfolio, we characterize the asset returns that make these choices mutually consistent—namely, the equilibrium returns that support optimal behavior given the investor's opportunity set and the aggregate supply of assets—and then solve for their equilibrium risk premiums in the spirit of Merton's (1973) ICAPM. In equilibrium, each asset's expected excess return decomposes into three components: (i) a market risk premium compensating for exposure to systematic risk, (ii) a carbon externality premium proportional to the asset's emissions intensity, and (iii) a climate-hedge premium compensating for covariance with shocks to the carbon concentration process. Under the simplifying assumption of zero covariance between innovations in CO_2 emissions and stock returns, so that the hedging demand is zero, we show that firms that are greener (browner) than the market portfolio with the same systematic risk earn negative (positive) alphas relative to the standard CAPM. These return differentials arise endogenously from the same mechanism that determines the SCC. This prediction is consistent with empirical evidence showing that firms with higher emissions exhibit higher average returns in both the US (Bolton and Kacperczyk, 2021; Hsu et al., 2023) and across 77 developed and developing countries (Bolton and Kacperczyk, 2023).

The SCC, which is a central metric in the literature on the economics of climate change (Golosov et al., 2014; Jensen and Traeger, 2021; van den Bremer and van der Ploeg, 2021), measures the expected present value of the long-run costs associated with emitting one unit of carbon today. In that literature, a policy maker sets a carbon tax equal to the SCC so that firms internalize the economic damage of their carbon emissions. In our setting, this mechanism is mirrored in financial markets: higher emissions translate into higher required returns because the representative investor internalizes the environmental damage of her investments. Firms that cause more damage therefore face higher capital costs. Importantly, our framework provides a measure of the SCC under risk, directly linked to information in financial markets. It captures the shadow price of carbon implied by investors' optimal behavior and can thus be inferred from observable asset prices.

We further derive an adjusted Sharpe ratio that incorporates climate concerns by adjusting for the SCC. Maximizing this *carbon-adjusted Sharpe ratio* is equivalent to solving the investor's original dynamic problem, and it provides a practical rule for portfolio selection. This micro-founded metric directly connects normative climate economics to

investment management with many risky assets and can be implemented using standard portfolio optimization tools.

To examine the mechanisms of the model in greater detail, we provide a closed-form solution under the assumptions of market completeness and a linear damage function. This solution yields analytical expressions for the SCC, as well as for optimal consumption and portfolio allocations. The characterization is valid for any value of the investor's elasticity of intertemporal substitution and thus nests power utility as a special case. The closed-form expressions allow us to quantify the impact of environmental damage on consumption and portfolio decisions in terms of model primitives. The presence of environmental damage reduces consumption and the myopic demand for carbon-intensive assets. From a hedging perspective, however, it might be optimal to hold more of the polluting asset if its return is positively correlated with CO_2 emissions, as such firms hedge against declines in non-financial income.

To illustrate the model's quantitative implications, we take it to US stock market data spanning 2007-2019. We estimate an implied SCC of about 53 USD per ton of carbon, consistent with benchmark values in the climate-economics literature. The SCC is stable across parameterizations, increasing with the damage elasticity and climate volatility but declining with the discount rate and the carbon absorption rate. Climate externalities modestly alter portfolio composition while leaving overall risk exposure intact: investors tilt slightly toward less carbon-intensive industries, and even under high-SCC scenarios weight shifts across industries or carbon-intensity quartiles remain below 5–6%. Hedging motives are negligible, as weak correlations between carbon dynamics and asset returns render the hedge component unimportant. The implied impact on equilibrium risk premiums is also modest, not exceeding 147 basis points per year between the most and least carbon-intensive firms (and 106 basis points across industries). In interpreting these results, it is important to note that the portfolio allocations we report are normative in that they describe how a climate-conscious investor would allocate wealth given prevailing asset prices and the SCC implied by the model. Equilibrium returns, by contrast, arise only once market clearing is imposed and the representative investor holds the aggregate market portfolio. The normative allocations therefore illuminate the mechanisms through which climate damages influence investment decisions, even though the market portfolio itself is determined in equilibrium.

Moving beyond the analytical benchmark, we solve the Merton–Nordhaus model numerically using a finite-difference approximation method to allow for market incompleteness, nonlinear damages, and binding short-sale constraints. The results confirm that the closed-form solution remains remarkably accurate. Short-sale constraints matter more than climate risk itself: when short sales are disallowed, climate-aware portfolios become more concentrated, as investors cannot offset carbon exposure through negative positions. This mechanical constraint amplifies cross-asset differences, not necessarily

because of strong climate premiums but due to limited diversification capacity. Despite these relaxations, the closed-form solution closely tracks the numerical results, underscoring that the key mechanisms are well captured by the analytical model and that nonlinearities play only a limited quantitative role.

Related literature. Our work contributes to a growing literature on climate and investments by offering a damage-based, long-horizon equilibrium model with uncertainty about climate impacts. The literature can be divided into two strands: utility-based models, which introduce explicit environmental preferences (Baker et al., 2022; Hsu et al., 2023; Pastor et al., 2021; Pedersen et al., 2021; Sauzet and Zerbib, 2022), and damage-based models, which incorporate climate effects through a damage function affecting production (Barnett, 2023; Barnett et al., 2020; Dietz et al., 2018; Hambel et al., 2024; Hong et al., 2023; Karydas and Xepapadeas, 2022; van den Bremer and van der Ploeg, 2021).

Utility-based models typically represent environmental preferences either at the governmental level (Hsu et al., 2023), or through heterogeneous individual investors (Baker et al., 2022; Pastor et al., 2021; Pedersen et al., 2021), sometimes with noisy ESG proxies (Avramov et al., 2022). These preferences influence asset prices and may also shape corporate environmental policies (Heinkel et al., 2001). However, they do not address climate change directly, but only indirectly through broader notions of environmental costs, sustainability, and/or ESG factors (environmental, social, and governance considerations).

By contrast, damage-based models follow the IAM tradition, using standard preferences for consumption and where climate affects the economy via an endogenous damage function that translates global warming into economic losses. Bansal et al. (2016) consider a social planner who jointly determines consumption and abatement and explore the link between equilibrium returns and the SCC, but without solving the portfolio choice problem. Karydas and Xepapadeas (2022) study rare climate disasters and their impact on the equity premium. In a two-firm setting, Hambel et al. (2024) show that diversification and climate goals may align in the short run but conflict in the long run. In an extension, Hambel and van der Ploeg (2025) show that sudden policy changes associated with an introduction of a carbon tax have important dynamic asset pricing consequences. Hong et al. (2023) analyze socially optimal investment mandates as an alternative to a carbon tax in order to meet zero net carbon emissions. Following a similar reasoning as formally developed in our paper, Pedersen (2025) argues that higher cost of capital that captures the SCC could be a substitute for too low or too little carbon taxes implemented by policy makers. Empirical tests, however, show that the return premiums currently charged by the market are largely insufficient relative to standard estimates of the SCC (Chittaro et al., 2025). Going beyond the existing literature, which focuses on a simple dichotomy of green vs. brown firms, our approach allows for a generic number of firms capturing large observed heterogeneity across firms. For example, although Chittaro et al. (2025) show that carbon premiums will matter little for most firms, they can be very substantial

in the right tail of emission intensity (mostly brown energy production).

In line with these papers, we take the perspective of a passive investor who *exits* from emissions-intensive firms by divesting. Relatedly, Chen et al. (2025) examine the effects of divestment strategies in a setting with heterogeneous investors focusing on the market effects of exclusion-based (rule-of-thumb) divestment strategies. In contrast, we consider a representative investor who internalizes the negative externality through the logic of the SCC in a setting with rich asset heterogeneity, both in financial performance and carbon intensity, guided by detailed evidence from the US stock market. While our approach centers on divestment, an alternative mechanism often emphasized in the literature is the *voice* argument, whereby investors actively engage with firms and their shareholder rights to influence corporate behavior (Broccardo et al., 2022; Gollier and Pouget, 2022). Our model abstracts from such active ownership and instead focuses on how investors internalize carbon risk through their portfolio decisions.

While the classic Dynamic Integrated Climate-Economy (DICE) model of Nordhaus (1993) abstracts from uncertainty, later studies show that stochastic elements in the climate or the economic process can substantially increase the SCC (Barnett et al., 2020; Hambel et al., 2021a; van den Bremer and van der Ploeg, 2021). Dietz et al. (2018) incorporate stochastic uncertainty into integrated assessment models and show how climate-related risks affect risk-adjusted discounting through the concept of a climate beta. Because solving these models numerically is complex (Cai and Lontzek, 2019), we initially follow Golosov et al. (2014) in providing analytical expressions within a tractable setup and under some simplifying assumptions —an approach that delivers transparent insights at the cost of some simplifications of geophysical dynamics and the damage functions.

Our model's prediction that carbon-intensive assets command higher required returns is broadly consistent with the evidence in Bolton and Kacperczyk (2021, 2023). At the same time, the small magnitude of carbon damages to equilibrium risk premiums and the economically insignificant carbon hedge help explain why the empirical literature finds mixed results on realized returns (Aswani et al., 2023; Zhang, 2025): while Bolton and Kacperczyk (2021) report higher required returns for brown firms, Pastor et al. (2022) and Bauer et al. (2022) document higher realized returns for green firms. In our framework, these patterns can coexist because the carbon-risk component of returns is small relative to the overall market component.

Our findings also speak to the debate on the market impact of sustainable investing. De Angelis et al. (2023) argue that a low share of green investors implies limited price effects. We show that even if all investors are climate-conscious, and under short-sale constraints tilt away from high-intensity assets, the resulting shifts in required returns remain modest under empirically plausible parameters.

The remainder of the paper is organized as follows. Section 2 develops the model and presents the main theoretical results, including a closed-form solution under a specific set

of assumptions. Section 3 documents empirical patterns in carbon intensities for a large sample of US firms. Section 4 quantifies the model's implications, estimates the SCC, analyzes optimal portfolios, and decomposes asset returns. Section 5 concludes.

2. A Merton-Nordhaus model

This section develops a unified Merton–Nordhaus model that combines Merton's dynamic portfolio framework (Merton, 1971) with Nordhaus-style damage-based IAM (Nordhaus, 1993). We follow the reduced-form treatment of climate damages in Golosov et al. (2014) to retain analytical tractability. In addition, Merton's (1973) ICAPM provides a dynamic equilibrium link between optimal portfolio choice and state variables. In our context, it is useful because it accommodates time-varying premiums and covariation between asset returns and climate risks, allowing us to study how exposure to changes in atmospheric CO_2 affects both optimal demands and equilibrium expected returns.

2.1 Model setup

We consider a representative investor with an infinite planning horizon whose investment opportunity set comprises N risky financial assets, and a riskless asset (a short-term bond). She optimizes her objective by choosing a consumption path $\{C(t)\}_{t=0}^{\infty}$ and portfolio weights $\{\mathbf{w}(t)\}_{t=0}^{\infty}$, where $\mathbf{w}(t) = [w_1(t), w_2(t), \cdots, w_N(t)]^\top$ denotes the fraction of financial wealth invested in each risky asset. Let $\mathbf{1}$ denote an $N \times 1$ vector of ones. The portfolio weight on the riskless asset is then given by $w_f(t) = \left(1 - \mathbf{w}(t)^\top \mathbf{1}\right)$.

Prices of risky financial assets $i = 1, 2, ..., N$ follow stochastic processes of the form

$$\frac{\mathrm{d}S_i(t)}{S_i(t)} = R_i \mathrm{d}t + \boldsymbol{\gamma}_i^\top \mathrm{d}\mathbf{Z}(t), \tag{2.1}$$

where $\mathbf{Z}(t)$ is an $N \times 1$ vector of independent Brownian motions and $\boldsymbol{\gamma}_i = [\gamma_{i1}, \gamma_{i2}, ..., \gamma_{iN}]^\top$ is an $N \times 1$ vector of constant diffusion loadings. Further, we define the $N \times 1$ vector of instantaneously expected returns (also assumed to be constant) as $\mathbf{R} = [R_1, R_2, \cdots, R_N]^\top$ and the $N \times N$ matrix of diffusion loadings $\boldsymbol{\Omega}^{1/2} = \left[\boldsymbol{\gamma}_1^\top, \boldsymbol{\gamma}_2^\top, \cdots, \boldsymbol{\gamma}_N^\top\right]^\top$. The instantaneous covariance matrix of returns is thus given by $\boldsymbol{\Omega} = \boldsymbol{\Omega}^{1/2}\left(\boldsymbol{\Omega}^{1/2}\right)^\top$. Finally, the risk-free bond is assumed to pay a constant rate of return, $r > 0$.

We assume that the representative investor has recursive preferences as initially proposed by Kreps and Porteus (1978), Epstein and Zin (1989) and Weil (1990), and extended to continuous time by Duffie and Epstein (1992a,b). This form of non-additive preferences has become popular in asset pricing and macroeconomics for at least two reasons. First, it allows us to disentangle the effects that risk aversion and intertemporal substitution have separately on optimal investment and consumption decisions. Secondly, it provides

a better quantitative and empirical fit for various asset pricing facts than the classical constant relative risk aversion (CRRA) utility, as demonstrated by Bansal and Yaron (2004) and the subsequent long-run risk literature. More specifically, the preferences of the investor, modulo a transversality condition, are given by

$$V(t) = \mathbb{E}_t \left[\int_t^\infty f(C(s), V(s)) \mathrm{d}s \right], \tag{2.2}$$

where $f(C, V)$ is a normalized aggregator of the current consumption rate, C, and of an appropriately transformed utility process, V, and $\mathbb{E}_t[\cdot]$ is the expectation operator conditional on the information available at time t. In its more general form, the aggregator is defined as

$$f(C, V) = \frac{1}{1 - \frac{1}{\epsilon}} C^{1-\frac{1}{\epsilon}} \left((1-\gamma) \rho V \right)^{1-\frac{1}{\zeta}} - \rho \zeta V, \tag{2.3}$$

where $\rho > 0$ is the rate of time preference, $\gamma > 0$ denotes the coefficient of relative risk aversion (RRA) towards atemporal bets, $\epsilon > 0$ denotes the elasticity of intertemporal substitution (EIS), so that $1/\epsilon$ can be understood as aversion towards intertemporal fluctuations, and $\zeta := (1-\gamma)/(1-\epsilon^{-1})$. The normalized aggregator exhibits the property that for $\gamma > 1/\epsilon$, the investor prefers early over later resolution of uncertainty.[3] The standard power CRRA utility used in Nordhaus (1993) can be recovered from (2.3) whenever $\gamma\epsilon = 1$ so $\zeta = 1$. In this case, $f(C, V) = \frac{C^{1-\gamma}}{1-\gamma} - \rho V$ which is time-additively separable in C and V, and further becomes the logarithmic utility if $\epsilon = \gamma = 1$.

Given the asset price dynamics, the representative investor's financial wealth evolves according to

$$\mathrm{d}W(t) = \left(rW(t) + \mathbf{w}(t)^\top (\mathbf{R} - r\mathbf{1}) W(t) + Y(t) - C(t) \right) \mathrm{d}t + \mathbf{w}(t)^\top \mathbf{\Omega}^{1/2} W(t) \mathrm{d}\mathbf{Z}(t), \tag{2.4}$$

with $Y(t)$ denoting the flow of non-financial income. This term is crucial in our model as it captures background risk from climate change. One can for instance think of non-financial income as income from other factors such as labor or land. It is well documented that climate change will depress labor productivity (Huang et al., 2020). Moreover, it will have adverse effects on productivity of agricultural land and also be associated with destruction of physical properties such as real estate (Giglio et al., 2021). The increasing prevalence of natural disasters, such as e.g. floodings or wild fires associated with climate change, will therefore negatively affect this income stream. Note that while in the model climate change does not have a direct impact on financial returns, these are nonetheless correlated with the state of the climate, as will become evident below.

The portfolio selection problem is paired with an environmental externality. Each financial investment in a specific firm i, $w_i(t)W(t)$, translates into output of CO_2 by means

[3] For the case of unitary EIS, $\epsilon = 1$, (2.3) reduces to $f(C, V) = \rho(1-\gamma)J \left(\ln C - \frac{1}{1-\gamma} \ln\left((1-\gamma)\rho J \right) \right)$.

of a linear mapping which simply entails a multiplication by a constant $k_i \geq 0$.[4] This constant is specific to each firm. The larger the value of k_i, the more firm i contributes to the world's concentration of CO_2. We collect the constants for all firms into the $N \times 1$ vector $\mathbf{k} = [k_1, \ldots, k_N]^\top$. Given that we consider a representative investor, her decision will determine the overall flow of anthropogenic CO_2 emissions

$$E(t) = \mathbf{k}^\top \mathbf{w}(t) W(t). \tag{2.5}$$

Thus, in our model, climate change is not captured by an exogenous process, but it is shaped by the representative investor's decisions. Throughout, we assume that the riskless asset has no carbon impacts, i.e., $k_f = 0$. Equation (2.5) mirrors the linear emissions mapping used in Nordhaus-style IAMs where emissions equal carbon intensity times aggregate output. It also parallels the logic of the Kaya identity where total emissions arise as the product of an economic scale term, financial wealth $W(t)$, and a portfolio-determined carbon intensity, $\mathbf{k}^\top \mathbf{w}(t)$, which plays the role of emissions per unit of economic activity.[5]

Besides the representative investor's wealth, we have to account for the stock concentration of CO_2 in the world. With the scalar variable η we capture the cumulative emissions of CO_2 in the atmosphere in excess of pre-industrial ages. It evolves according to the following relation:

$$\mathrm{d}\eta(t) = (E(t) - \delta\eta(t))\,\mathrm{d}t + \sigma_\eta \eta(t) \left(\hat{\boldsymbol{\rho}}_{\eta S}^\top \mathrm{d}\mathbf{Z}(t) + \sqrt{1 - \|\hat{\boldsymbol{\rho}}_{\eta S}\|^2}\, \mathrm{d}B(t) \right), \tag{2.6}$$

where $B(t)$ is a one-dimensional Brownian motion independent of the Brownian motions in $\mathbf{Z}(t)$. The process $B(t)$ captures unexpected variations in the concentration of CO_2 not stemming from anthropogenic emissions such as, e.g., volcano eruptions. On the other hand, the process $\mathbf{Z}(t)$ captures components that affect both the price dynamics of risky assets and the stock of CO_2. For example, (unanticipated) legal restrictions on emissions will impact both stock of CO_2 as well as the returns of carbon-intense businesses such as airlines. That is, the process $\mathbf{Z}(t)$ includes transition risk.

The constant σ_η denotes the volatility of the growth in cumulative emissions of CO_2, and the vector $\hat{\boldsymbol{\rho}}_{\eta S} = [\hat{\rho}_{\eta 1}, \ldots, \hat{\rho}_{\eta N}]^\top$ captures the instantaneous correlation between

[4] Theoretically, this value could even be negative if a firm reduces the stock of carbon in the atmosphere, e.g., by planting trees that act as a sink for the carbon concentration. Given current industry practices this seems like an unrealistic assumption. Technically, together with Equation (2.6) this could also imply a negative stock of carbon in the atmosphere η, which is clearly unrealistic.

[5] The Kaya identity decomposes aggregate CO_2 emissions as the product of an economic scale component and an emissions-intensity component. Our expression $E(t) = \mathbf{k}^\top \mathbf{w}(t)\, W(t)$ has an analogous structure: aggregate financial wealth $W(t)$ plays the role of the economic scale variable, while $\mathbf{k}^\top \mathbf{w}(t)$ represents the portfolio-weighted carbon intensity of production, paralleling the "emissions per unit of economic activity" terms in the Kaya identity. Unlike the classical identity, however, emissions intensity in our setting is endogenously determined by investors' portfolio decisions.

carbon concentration and asset returns.[6] In fact, these values will be of crucial importance for the hedging component. The structure of Equation (2.6) is such that it guarantees that $\eta(t) > 0$ almost surely (a.s.) at all points in time. Some of the carbon output is captured in the existing biosphere or the oceans, which we model as having a natural absorption rate of $\delta > 0$. In the absence of human activities ($E(t) = 0$) the carbon concentration will converge back to its pre-industrial level ($\eta = 0$). The stochastic ecological budget constraint in Equation (2.6) provides us with a connection running from financial activity to the environment.[7]

The reverse connection running from the ecological domain to the financial one is captured by means of a damage function. In particular, we assume that the representative investor's flow of non-financial income, Y, is affected by the state of CO_2 concentrations, η, through an environmental damage function $D(\eta)$, according to

$$Y(t) = Y(0)(1 - D(\eta(t))), \tag{2.7}$$

where $Y(0) = Y_0 > 0$ is the level of non-financial income at time 0. The relation in Equation (2.7) is consistent with the linear relationship used in IAMs, and in particular to that in Golosov et al. (2014), where the damage function captures the mapping from the carbon concentrations to economic damages measured as a fraction of total output.[8] A substantial body of literature focuses on the challenging task of quantifying the effects of climate change on economic output by assuming that the environmental damage function is a strictly increasing function, which takes on values between zero and one, i.e., $0 \leq D(\eta) \leq 1$ for all η (cf. Hassler et al., 2016). For example, Golosov et al. (2014) propose a convex mapping of the type

$$D(\eta(t)) = 1 - \exp(-\psi\eta(t)), \tag{2.8}$$

where $\psi > 0$ is a scaling parameter, the damage elasticity, that captures the adverse effects of carbon concentrations on economic activity. We view the above specification as a reduced-form way of linking climate dynamics to investors' wealth, whereby climate damages affect the investor through the non-financial income process, while expected

[6] More precisely, $\hat{\boldsymbol{\rho}}_{\eta S} = \boldsymbol{\Lambda}^{-1}\boldsymbol{\rho}_{\eta S}$, with $\boldsymbol{\Lambda}$ a Cholesky factorization of the correlation matrix of asset returns $\mathbf{C} = [\rho_{ij}]_{i,j=1,\ldots,N}$, and $\boldsymbol{\rho}_{\eta S}$ the correlation between between carbon concentration and asset returns. See Appendix A in the Online Supplementary Material for more details.

[7] We follow Golosov et al. (2014) in representing climate change through cumulative carbon emissions. A more common approach in both the economics and climate science literature is to model damages as a function of temperature increases relative to pre-industrial levels. However, as shown by Matthews et al. (2009) and applied in an IAM by Dietz and Venmans (2019), temperature changes are approximately proportional to cumulative emissions, rendering our formulation effectively isomorphic to the standard temperature-based one.

[8] An alternative approach to capture the feedback from the environment to the economy is to include the climate externality directly into the utility function of investors (see Acemoglu et al., 2012 and Barrage, 2020).

asset returns and return covariances, $\mathbf{R}$ and $\mathbf{\Omega}$, are treated as exogenous primitives of the investment opportunity set in the spirit of Merton (1971). This separation preserves analytical tractability and allows us to isolate how the Social Cost of Carbon and climate risk enter portfolio choice and equilibrium pricing.[9]

In summary, the problem faced by the representative investor is that of maximizing the recursive utility given in (2.2)–(2.3) subject to the wealth accumulation process (2.4), and the process for the stock of CO_2 concentration (2.5)-(2.6). Applying the results in Duffie and Epstein (1992a,b), we find that combining these model elements gives rise to the following Hamilton-Jacobi-Bellman (HJB) equation:

$$\begin{aligned} 0 = \max_{C,\mathbf{w}} \Big\{ & f(C,J) + \left[rW + \mathbf{w}^\top(\mathbf{R} - r\mathbf{1})W - Y_0 D(\eta) + (Y_0 - C) \right] J_W \\ & + \left[\mathbf{k}^\top \mathbf{w} W - \delta\eta \right] J_\eta + \frac{1}{2} J_{WW} \mathbf{w}^\top \mathbf{\Omega} \mathbf{w} W^2 + J_{W\eta} \mathbf{w}^\top \mathbf{\Omega}^{1/2} \sigma_\eta \eta \hat{\boldsymbol{\rho}}_{\eta S} W + \frac{1}{2} J_{\eta\eta} \sigma_\eta^2 \eta^2 \Big\}, \end{aligned} \tag{2.9}$$

where $J := J(W,\eta)$ denotes the value function, representing the investor's maximum utility given financial wealth W and the CO_2 concentration η. Subscripts denote partial derivatives.[10]

2.2 Individual portfolio optimization

To identify optimal consumption and portfolio weights, we solve the problem in Equation (2.9). The following proposition summarizes the representative investor's optimal rules.

Proposition 1 (Optimal consumption and portfolio weights)**.** *The rate of optimal consumption, $C := C(W,\eta)$, in the Merton-Nordhaus model follows from*

$$f_C(C,J) = J_W(W,\eta). \tag{2.10}$$

The optimal portfolio weights, $\mathbf{w} := \mathbf{w}(W,\eta)$, are given by

$$\mathbf{w} = \mathbf{\Omega}^{-1}(\mathbf{R} - r\mathbf{1})\theta + \mathbf{\Omega}^{-1}\mathbf{k}\kappa + \mathbf{\Omega}^{-1}\mathbf{\Omega}^{1/2}\sigma_\eta \eta \hat{\boldsymbol{\rho}}_{\eta S} H, \tag{2.11}$$

where

$$\theta := \theta(W,\eta) = -\frac{J_W(W,\eta)}{J_{WW}(W,\eta)W}, \tag{2.12}$$

[9] Allowing expected asset returns ain Equation (2.1) to depend directly on the climate state $\mathbf{R}(\eta)$, for instance through climate-sensitive cash-flow growth, would not change the qualitative structure of our theoretical results. It would mainly scale the components without changing their meaning, but at the cost of losing tractability and the closed-form characterization. For clarity, we therefore maintain a constant investment opportunity set.

[10] Note the distinction between the value function J and the transformed utility process V in Equation (2.2): V represents utility for any admissible consumption path, whereas J corresponds to the maximum given optimal controls.

$$\kappa := \kappa(W, \eta) = -\frac{J_\eta(W, \eta)}{J_{WW}(W, \eta)W}, \tag{2.13}$$

and

$$H := H(W, \eta) = -\frac{J_{W\eta}(W, \eta)}{J_{WW}(W, \eta)W}. \tag{2.14}$$

Proof. The result follows from the first-order conditions to the HJB equation in (2.9). ■

From Equation (2.10), the optimal level of consumption is such that the marginal benefit of an additional unit of consumption is equal to the marginal utility of an additional unit of financial wealth. Equation (2.11) determines the optimal share of financial wealth allocated to the risky assets as the sum of three components. The first one captures the myopic demand (Merton, 1969, 1971). The second term is related to the negative externality arising from pollution, which in turn lowers future consumption. In the former, assets are weighted by their expected excess returns, whereas in the latter, weights associated with the individual assets are related to the technological parameter $\mathbf{k}$ capturing the CO_2 output per dollar invested. Finally, the third term constitutes an intertemporal hedging demand or excess risky demand (Merton, 1971), which hedges against unexpected, unfavorable changes in the amount of carbon in the atmosphere.[11]

Proposition 1 implies a four-fund separation result, according to which the representative investor is indifferent between investing in the full menu of assets and investing in four funds, namely the risk-free asset, a myopic fund, a climate-externality fund, and a climate-hedge fund. Here, the myopic fund corresponds to the usual tangency portfolio emerging from mean-variance analysis and the hedge fund to that in Merton (1973).

Proposition 2 (Four-fund separation)**.** *The representative investor is indifferent between investing in the $N + 1$ assets in the investment opportunity set and investing in four funds: 1) a risk-free asset; 2) a portfolio of risky assets with weights*

$$\mathbf{w}_{\text{myopic}} = \frac{\mathbf{\Omega}^{-1}(\mathbf{R} - r\mathbf{1})}{\mathbf{1}^\top \mathbf{\Omega}^{-1}(\mathbf{R} - r\mathbf{1})}; \tag{2.15}$$

3) a portfolio of risky assets with weights

$$\mathbf{w}_{\text{ext}} = \frac{\mathbf{\Omega}^{-1}\mathbf{k}}{\mathbf{1}^\top \mathbf{\Omega}^{-1}\mathbf{k}}; \tag{2.16}$$

[11] This hedging component differs from the standard intertemporal hedge in ICAPM models. In the usual setting, hedge demands arise from shocks to the investment opportunity set and operate through cross-asset return covariances in the myopic demand via the covariance matrix $\mathbf{\Omega}$. Here, instead, hedging is tied to the climate state: an asset hedges only to the extent that its returns covary with shocks to atmospheric carbon η. Cross-asset correlations still matter for portfolio choice through $\mathbf{\Omega}$, but they are separate from, and conceptually distinct from, the climate-hedging channel emphasized here.

and 4) a portfolio of risky assets with weights

$$\mathbf{w}_{\text{hedge}} = \frac{\boldsymbol{\Omega}^{-1}\boldsymbol{\Omega}^{1/2}\hat{\boldsymbol{\rho}}_{\eta S}}{\mathbf{1}^{\top}\boldsymbol{\Omega}^{-1}\boldsymbol{\Omega}^{1/2}\hat{\boldsymbol{\rho}}_{\eta S}}. \tag{2.17}$$

Proof. The result follows immediately from Proposition 1. The weights on each of the funds are $\alpha_{\text{myopic}} = \theta\,\mathbf{1}^{\top}\boldsymbol{\Omega}^{-1}(\mathbf{R} - r\mathbf{1})$, $\alpha_{\text{ext}} = \kappa\,\mathbf{1}^{\top}\boldsymbol{\Omega}^{-1}\mathbf{k}$, $\alpha_{\text{hedge}} = H\,\mathbf{1}^{\top}\boldsymbol{\Omega}^{-1}\boldsymbol{\Omega}^{1/2}\sigma_{\eta}\eta\hat{\boldsymbol{\rho}}_{\eta S}$, and $\alpha_{\text{risk-free}} = 1 - \alpha_{\text{myopic}} - \alpha_{\text{ext}} - \alpha_{\text{hedge}}$. ■

In our setup, rational investors account for the climate externality induced by the private firms. Not only can they support climate-friendly companies by overweighting them in the portfolio, but also punish polluting companies by shorting their assets. In lieu of a carbon tax, which is a direct cost for the firm, here the costs of carbon are reflected in higher financing costs for the firm imposed by the investors. The novel hedge term, in turn, captures the background risk induced by climate change. As the representative investor internalizes these effects, they are ultimately reflected in equilibrium returns, which we analyze in the following.

2.3 Return decomposition

Similar to Merton (1973), we use the representative investor's optimal asset demands to derive an equilibrium return decomposition, i.e., a relation consistent with optimal portfolio and consumption choices, though not derived from a general-equilibrium setting with endogenous prices.[12] Modulo a technical invertibility condition (given below), the equilibrium risk premiums of risky assets 1 through $(N-2)$ can be written as linear combinations of the market risk premium and the risk premiums of assets $(N-1)$ and N, respectively. Of course, the ordering of risky assets is arbitrary, and one can create (redundant) assets as linear combinations of the original assets, so it suffices that there are two assets or portfolios of assets that satisfy the invertibility condition.

Proposition 3 (Equilibrium return decomposition)**.** *Suppose the* $(N-1)$*th asset's return has an instantaneous covariance of* $\sigma_{(N-1)\eta}$ *with changes in* η*, while the* N*th asset's returns has an instantaneous covariance of* $\sigma_{N\eta}$ *with changes in* η*. We also define the* 3×3 *matrix*

$$\boldsymbol{\Xi} = \begin{bmatrix} \sigma_M^2 & k_M & \sigma_{M\eta} \\ \sigma_{M(N-1)} & k_{N-1} & \sigma_{(N-1)\eta} \\ \sigma_{MN} & k_N & \sigma_{N\eta} \end{bmatrix}, \tag{2.18}$$

[12] We derive an ICAPM equilibrium where the investment opportunity set is exogenous and the market portfolio weights are endogenously determined. While the ICAPM is sometimes described as a general equilibrium model, we follow the interpretation of a partial equilibrium framework in which the representative investor optimizes consumption and portfolio choices taking prices as given. Under linear production technologies (see Cox et al., 1985a,b), however, the same structure can be embedded in a general-equilibrium environment with endogenous asset prices; see Cochrane (2021) for discussion.

where M stands for the market portfolio and $k_M := \mathbf{w}_M^\top \mathbf{k}$ (and the rest of the notation is self-explanatory). Provided that $\boldsymbol{\Xi}^{-1}$ exists, the expected return of assets $i = 1, 2, ..., (N-2)$ can be expressed as

$$R_i - r = \begin{bmatrix} \sigma_{iM} & k_i & \sigma_{i\eta} \end{bmatrix} \boldsymbol{\Xi}^{-1} \begin{bmatrix} R_M - r \\ R_{N-1} - r \\ R_N - r \end{bmatrix} = \begin{bmatrix} \sigma_{iM} & k_i & \sigma_{i\eta} \end{bmatrix} \begin{bmatrix} \frac{1}{\theta} \\ -\frac{\kappa}{\theta} \\ -\frac{H}{\theta} \end{bmatrix}. \tag{2.19}$$

Proof. See Appendix A. ■

Equation (2.19) provides us with intuition regarding the drivers of expected excess returns by decomposing the risk premium into three components, each corresponding to a specific aspect of the portfolio of risky assets in Proposition 1: (i) the market risk premium, (ii) the carbon externality premium, and (iii) the climate-hedge premium. The market risk premium captures exposure to aggregate market fluctuations. Assets that co-move strongly with the market must offer higher returns to compensate for bearing systematic market risk. This is the standard CAPM component. The carbon externality premium reflects exposure to firms' carbon footprints. Under the assumption that the investor's indirect utility is increasing in wealth ($J_W > 0$) and decreasing in the stock of atmospheric CO_2 emissions ($J_\eta < 0$), we have $-\kappa/\theta > 0$. This implies that the premium is positive for firms that contribute to higher carbon emissions: "dirty" firms must offer higher expected returns, otherwise investors would not hold them in equilibrium. Finally, the climate-hedge premium captures exposure to shocks in accumulated CO_2 concentrations. If the marginal utility of wealth rises with pollution ($J_{W\eta} > 0$),[13] then $-H/\theta < 0$. Consequently, this premium is negative for stocks whose returns are positively correlated with increases in η, and positive for stocks whose returns are negatively correlated. Intuitively, assets that perform well when climate risks intensify act as hedges; they are in higher demand and can therefore offer lower expected returns in equilibrium.

Note that the risk premium associated with "a unitary contribution to η" (the case when $k_i = 1$) is given by[14]

$$R_i - r\Big|_{k_i=1} = -\frac{\kappa}{\theta} = -\frac{J_\eta(W,\eta)}{J_W(W,\eta)} = -\frac{J_\eta(W,\eta)}{f_C(C,J)} := \text{SCC} > 0, \tag{2.20}$$

which defines our measure of the *Social Cost of Carbon* (SCC), i.e., the shadow price of carbon implied by optimal investor behavior. It represents the additional cost of capital required by a rational representative investor with an infinite planning horizon who internalizes the economic damage caused by her investment decisions. Our metric is equivalent to that employed in the literature on the economics of climate change where

[13] The intuition is that since η is a "bad," the marginal utility of wealth increases with η, i.e. $J_{W\eta} > 0$.

[14] The second equality uses Equations (2.12) and (2.13), and the third equality follows from the envelope condition in Equation (2.10).

the SCC, or market price of carbon, is the carbon tax that a benevolent social planner would set in order to make firms internalize the economic damage caused by their carbon emissions. It also coincides with the definition used in van den Bremer and van der Ploeg (2021), where the SCC is the marginal rate of substitution between carbon emissions and economic activity.

To illustrate the implications of the decomposition in Proposition 3, consider a scenario in which the stock of CO_2 in the atmosphere is uncorrelated with the traded assets, so that climate risk cannot be hedged through financial markets. Suppose further that the Nth asset has a zero beta with the market ($\sigma_{MN} = 0$) and induces CO_2 emissions ($k_N > 0$). Under these assumptions, the risk premium of asset $i = 1, 2, ..., N-1$ is given by

$$R_i - r = \frac{1}{\theta}\sigma_{iM} - \frac{\kappa}{\theta}k_i. \tag{2.21}$$

Further, we get the following system of equations, based on the risk premium of the market and that of asset N, respectively:

$$R_M - r = \frac{1}{\theta}\sigma_M^2 - \frac{\kappa}{\theta}k_M, \tag{2.22}$$

$$R_N - r = -\frac{\kappa}{\theta}k_N. \tag{2.23}$$

Since $k_N > 0$ and $-\frac{\kappa}{\theta} > 0$, it follows that the risk premium of asset N is positive, i.e., $R_N - r > 0$. Solving for $\frac{1}{\theta}$ and $-\frac{\kappa}{\theta}$ in the system formed by Equations (2.22) and (2.23), and inserting the expressions into Equation (2.21), we arrive at

$$R_i - r = \left(\frac{k_i - \beta_i k_M}{k_N}\right)(R_N - r) + \beta_i(R_M - r), \tag{2.24}$$

where $\beta_i = \sigma_{iM}/\sigma_M^2$. Hence, if an asset induces more (less) emissions than what is suggested by its market exposure (k_i compared to $\beta_i k_M$), it will display a positive (negative) alpha relative to the traditional CAPM, given $R_N - r > 0$. In other words, heavy polluters will appear to generate positive alphas, while environmentally-friendly firms will appear to generate negative alphas. This result is in line with empirical findings that firms with higher carbon emissions exhibit higher returns (Bolton and Kacperczyk, 2021, 2023).

With the SCC identified from Equation (2.20), it is useful to revisit the investor's optimal allocation problem. In particular, if the hedging demand is zero, the original problem in Section 2.2 simplifies to an adjusted mean–variance formulation, or equivalently, to the maximization of a Sharpe ratio corrected for the SCC, as shown in the next proposition.

Proposition 4 (Carbon-adjusted Sharpe ratio)**.** *Suppose that the hedging demand is zero. Then, the investor's portfolio optimization problem is equivalent to the following adjusted*

mean-variance optimization problem:

$$\max_{\mathbf{w}} \mathbf{w}^\top (\mathbf{R} - \text{SCC}\,\mathbf{k}) + \left(1 - \mathbf{w}^\top \mathbf{1}\right) r - \frac{1}{2}\theta^{-1}\mathbf{w}^\top \mathbf{\Omega}\mathbf{w}, \tag{2.25}$$

where $\frac{1}{\theta} = -\frac{WJ_{WW}}{J_W}$ *is a measure of relative risk aversion (with respect to the value function). Further, assume that* $\mathbf{1}^\top \mathbf{\Omega}^{-1}(\mathbf{R} - \text{SCC}\,\mathbf{k} - r\mathbf{1}) > 0$*. Then, maximizing the adjusted Sharpe ratio of a risky portfolio with portfolio weights* $\widehat{\mathbf{w}}$*,*

$$\begin{aligned} \max_{\widehat{\mathbf{w}}} \quad & \frac{\widehat{\mathbf{w}}^\top (\mathbf{R} - \text{SCC}\,\mathbf{k}) - r}{\sqrt{\widehat{\mathbf{w}}^\top \mathbf{\Omega}\widehat{\mathbf{w}}}} \\ s.t. \quad & \widehat{\mathbf{w}}^\top \mathbf{1} = 1, \end{aligned} \tag{2.26}$$

results in an adjusted mean-variance optimal risky asset portfolio. The solution to the adjusted mean-variance optimal portfolio problem in (2.25) can be found by combining the risk-free asset and the risky asset portfolio that solves the optimization problem in (2.26).

Proof. See Appendix B.1 in the Online Supplementary Material. ■

The interpretation is that if conditions are met for a myopic portfolio choice problem, then the solution to the dynamic portfolio choice problem can be found by considering an adjusted mean-variance problem in which one adjusts for the social cost of the portfolio's carbon emissions. Naturally then, with this adjustment, we obtain results that are very similar to the ones obtained in classic mean-variance portfolio theory.

The following corollary shows that Proposition 4 remains valid in the presence of short-selling constraints.

Corollary 1 (Carbon-adjusted Sharpe ratio under shortselling constraints)**.** *Provided that* $(\widehat{\mathbf{w}}^\star)^\top (\mathbf{R} - \text{SCC}\,\mathbf{k}) - r > 0$*, where* $\widehat{\mathbf{w}}^\star$ *is the maximizer of the adjusted Sharpe ratio, and that borrowing at the riskless rate is permitted, the result in Proposition 4 continues to hold under short-sale constraints* $\mathbf{w} \geq 0$*.*

Proof. See Appendix B.2 in the Online Supplementary Material. ■

2.4 A closed-form solution

In this section we provide an analytical solution to the consumption and asset allocation problem faced by the representative investor in the Merton-Nordhaus model. To obtain the solution we make the following two assumptions:

Assumption 1. *Markets are complete.*

Assumption 2. *The environmental damage function is linear in* CO_2 *concentrations.*

Assumption 1 requires that $\|\hat{\boldsymbol{\rho}}_{\eta S}\|^2 = 1$, so that climate risk, and hence income risk, is spanned by the financial market (see Munk and Sørensen, 2010). Thus, in the absence

of portfolio constraints, the dynamics for the concentration of CO_2 can be perfectly replicated by a dynamic trading strategy consisting of risky assets and the risk free rate. On the other hand, we implement Assumption 2 through a first-order Taylor approximation of the damage function in Equation (2.8). Specifically,

$$D(\eta(t)) \approx \psi\eta(t), \tag{2.27}$$

which is similar to the damage function used in Nordhaus (2015), and can be interpreted as a growth rate reduction, in line with Hambel et al. (2021b).

Under Assumptions 1 and 2, and provided the additional condition $\rho > (1-\epsilon^{-1})\hat{R}$ holds so the problem is well-posed, the following corollary to Proposition 1 summarizes the analytical functional form of the individual optimal rules.

Corollary 2 (Complete markets and linear damage)**.** *Let $X := W + \frac{Y_0}{r} - \mathrm{SCC}\,\eta$ denote the investor's total wealth. Then, optimal consumption is given by*

$$C = c_X X, \tag{2.28}$$

with $c_X = \hat{R} + \epsilon\,(\rho - \hat{R})$ the marginal propensity to consume and where the constant $\hat{R}$ is defined as

$$\hat{R} := r + \frac{1}{2}\frac{\Big((\mathbf{R} - r\mathbf{1})^\top - \mathrm{SCC}\,\mathbf{k}^\top\Big)\Big(\mathbf{\Omega}\Big)^{-1}\Big((\mathbf{R} - r\mathbf{1}) - \mathrm{SCC}\,\mathbf{k}\Big)}{\gamma}. \tag{2.29}$$

The optimal portfolio weights are given by

$$\mathbf{w} = \mathbf{\Omega}^{-1}(\mathbf{R} - r\mathbf{1})\theta + \mathbf{\Omega}^{-1}\mathbf{k}\kappa + \mathbf{\Omega}^{-1}\mathbf{\Omega}^{1/2}\sigma_\eta\eta\hat{\boldsymbol{\rho}}_{\eta S}H, \tag{2.30}$$

with

$$\theta = \frac{X}{\gamma W}, \quad \kappa = -\mathrm{SCC}\frac{X}{\gamma W}, \quad H = \mathrm{SCC}\frac{1}{W}, \tag{2.31}$$

where the social cost of carbon, $\mathrm{SCC} = -\frac{J_\eta}{f_C} = -\frac{J_\eta}{J_W}$, is constant and given by

$$\mathrm{SCC} = \left[\frac{1}{2a}\left(b - \sqrt{b^2 - 4ac}\right)\right] > 0, \tag{2.32}$$

with $a := \hat{\boldsymbol{\rho}}_{\eta S}^\top \left(\mathbf{\Omega}^{1/2}\right)^{-1}\mathbf{k}\sigma_\eta$, $b := \hat{\boldsymbol{\rho}}_{\eta S}^\top \left(\mathbf{\Omega}^{1/2}\right)^{-1}(\mathbf{R} - r\mathbf{1})\sigma_\eta + r + \delta$, and $c := Y_0\psi$. For the case of unitary EIS, $\epsilon = 1$, the marginal propensity to consume reduces to $c_X = \rho$, while the optimal portfolio weights and the SCC *remain the same.*

Proof. We guess that the value function takes the form $J(X) = c_X^{-\frac{\zeta}{\epsilon}}\rho^{\zeta-1}\frac{X^{1-\gamma}}{1-\gamma}$, with $X = W + \mathbb{C}_1 Y_0 + \mathbb{C}_2\eta$ for some unknown constants c_X, $\mathbb{C}_1$, and $\mathbb{C}_2$. Next, we verify if

the guess satisfies the HJB equation. See Appendix B.3 in the Online Supplementary Material for a detailed derivation. ■

The optimal consumption is given by the modified Keynes–Ramsey rule in Equation (2.28) whereby consumption is a linear function of investor's total wealth, X. The latter is the sum of financial wealth, W, and non-financial wealth measured as the present value of the non-financial income, Y_0/r, net of the environmental costs caused by the negative environmental externality from the optimal investment strategy, $\text{SCC}\,\eta$. The marginal propensity to consume out total wealth, c_X, is constant, and its value is determined by the EIS and the wedge between the subjective discount rate, ρ, and the constant, $\hat{R}$, i.e., $(\rho - \hat{R})$. We refer to the latter as the *risk-* and *carbon emissions-adjusted* return of the optimal portfolio (see van den Bremer et al., 2016). Note that the marginal propensity to consume also depends indirectly on the coefficient of relative risk aversion through this risk-adjusted return in Equation (2.29). For a given value of the EIS, the investor consumes more than her total wealth's return, $\hat{R}X$, whenever $\rho > \hat{R}$. Moreover, the response of consumption to changes in this wedge increases with larger values of the EIS. These results simplify in the case of an unitary EIS so the marginal propensity to consume is uniquely determined by the subjective discount rate. Overall, since $\text{SCC} > 0$, increased carbon concentrations lower the representative investor's optimal consumption.

Equation (2.30) provides the closed-form expressions for the different components of the optimal shares of financial wealth allocated to the risky assets. These are independent of the EIS (ϵ) as a result of the non-additivity assumption of the investor preferences. Besides the standard myopic portfolio demands trading off excess returns and volatility under some risk aversion γ, the environmental impact is also taken into account. Since $\text{SCC} > 0$ and $k_i \geq 0$, the myopic demand for risky assets of the representative investor is lower than the standard myopic demand without climate concerns since these assets also entail a negative externality for the environment. Ceteris paribus, the larger the carbon emission intensity of firm i, k_i, the lower is the fraction of financial wealth invested in that particular firm. Hence, the increasing prevalence of climate change will induce myopic individuals to hold less risky assets in general. This contractionary effect on the demand for risky assets is reinforced by the negative effect that the carbon concentration η has on total wealth. Specifically, as long as the carbon-intensity adjusted return is positive, i.e., $(\mathbf{R} - \text{SCC}\,\mathbf{k} - r\mathbf{1}) > 0$, we also observe a negative effect of carbon concentration on the myopic demand for risky assets running through X.

On the other hand, there can be a countervailing effect that runs through the intertemporal hedging demand for risky assets. In particular, investors will hold larger fractions of the risky assets —including brown assets with high carbon intensities— whenever their returns are highly correlated with the concentration of CO_2 emissions, i.e, when $\hat{\rho}_{\eta i}$ is sufficiently large for asset $i = 1, \ldots, N$. Suppose for example that the government (unexpectedly) reduces environmental regulations for companies. This increases the carbon

concentration and thereby induces a negative impact on non-financial income. At the same time, due to positive correlations, it tends to increase the returns of brown companies. Thus, by holding shares in brown companies, investors can hedge against the risk associated with environmental damages. This hedge term counteracts the lowered myopic demand originating from the climate externality. Note however that, as opposed to the standard case in Merton (1971), the hedging demand does not depend on the coefficient of relative risk aversion, γ. This is because, in our setting, the carbon concentration η affects non-financial income, but not the investment opportunity set.

The hedging term even plays a role for carbon-neutral companies, i.e., firms with $k_i = 0$. If their business model is negatively impacted by climate change so $\rho_{\eta i} < 0$, the representative investor will reduce the holdings in these companies. A practical example would be that of a reinsurance company: while its business itself does not contribute to carbon emissions, it bears the costs of climate change by e.g. offering flood insurance (Giglio et al., 2021).

Finally, the optimal value of the environmental externality, which in our framework equals the SCC, is also independent of the EIS (ϵ). Moreover, it depends in a non-trivial way on its discount rate.[15] To see this, it is insightful to consider some particular cases that are nested by Equation (2.32). For example, for $\sigma_\eta = 0$ the CO_2 concentrations evolve deterministically and hence[16]

$$\lim_{\sigma_\eta \to 0} \text{SCC} = \frac{Y_0 \psi}{r + \delta} > 0. \tag{2.33}$$

In this case the SCC, which we refer to as the *risk-free SCC*, crucially depends on three factors: (i) the absorption of carbon in the atmosphere δ (an ecologic factor), (ii) the risk-free rate r (an economic factor), and (iii) the damage elasticity ψ measuring the impact of carbon in the atmosphere on non-financial income connecting the economic and ecological domain of the model. In this case, the carbon concentration can be discounted at the risk-free rate. Moreover, the SCC decreases in δ, so if the absorption capacity of the ecosystem is high (high values on δ), the negative weight attached to the carbon concentration is lower in absolute value. As σ_η approaches zero, the hedging demand goes to zero and the optimal weights in Equation (2.30) reduce to

$$\mathbf{w} = \mathbf{\Omega}^{-1}(\mathbf{R} - r\mathbf{1} - \text{SCC}\ \mathbf{k})\frac{X}{\gamma W}.$$

Hence, the demands for the risky carbon-emitting assets increase with δ due to the negative effect of SCC on the numerator of said demands. Similar comparative statics to an

[15] Our model does not require specifying the rate of time preference ρ to compute the SCC, even though the SCC is typically sensitive to the chosen discount rate (Stern, 2007). We avoid this issue by taking the investment opportunity set as given.

[16] This assumption implies $a = 0$ and $b = r + \delta$ in Equation (2.32).

increase in δ carry through for a reduction in ψ, i.e, a less sensitive damage function.

Another interesting nested case emerges when $\psi = 0$, i.e., the case in which there is no feedback running from the climate to the economic sphere. In this case, the SCC $= 0$ which in turn implies that the value function does not depend on the stock of CO_2 concentrations, η. The interpretation of this result is straightforward: in the limiting case when the carbon concentration does not affect non-financial income, the investor does not care about it, and hence it does not enter the value function.

Next, consider the case in which $k_i = 0$ for all i so the carbon concentration evolves exogenously and is not influenced by the decisions of the representative investor. Here,

$$\text{SCC} = \frac{Y_0 \psi}{\hat{\boldsymbol{\rho}}_{\eta S}^{\top} \left(\boldsymbol{\Omega}^{1/2}\right)^{-1} (\mathbf{R} - r\mathbf{1})\sigma_\eta + r + \delta}. \tag{2.34}$$

Due to the spanning assumption, the risk associated with stochastic fluctuations in η can be totally hedged away. The correct discount rate for carbon concentrations is characterized by a premium (relative to the risk-free rate r) amounting to $\hat{\boldsymbol{\rho}}_{\eta S}^{\top} \left(\boldsymbol{\Omega}^{1/2}\right)^{-1} (\mathbf{R} - r\mathbf{1})\sigma_\eta$. This premium disappears if carbon evolves in a deterministic manner, i.e., if $\sigma_\eta = 0$.

Finally, one may also consider the case of an infinitely risk averse investor ($\gamma \to \infty$), or risk avoider, in which the optimal investment shares would neither display a myopic nor a carbon externality component. Nevertheless, she would expose herself to risky assets to hedge against the risk associated with the climate, since investing exclusively in the risk-free asset would not achieve this goal. The optimal portfolio weights in this case are

$$\mathbf{w} = \boldsymbol{\Omega}^{-1} \boldsymbol{\Omega}^{1/2} \hat{\boldsymbol{\rho}}_{\eta S} \sigma_\eta \frac{\text{SCC}\, \eta}{W}.$$

One might consider this as the diametric counterpart to the previous case in which carbon risk does not matter at all.

Given the general result regarding risk premiums that we proved in the previous section, we can easily derive the risk premiums in this specific setting.

Corollary 3 (Risk premiums)**.** *Let Assumptions 1 and 2 hold. Then, the risk premium of the ith risky asset is given by*

$$R_i - r = \sigma_{iM} \frac{\gamma W}{X} + k_i \,\text{SCC} - \sigma_{i\eta} \frac{\gamma \,\text{SCC}}{X}. \tag{2.35}$$

Proof. The result follows directly from the general asset pricing Equation (2.19). ■

Equation (2.35) captures two motives that are intertwined: (i) the relative risk aversion captured by $\gamma > 0$ (trading off risky vs. risk-free asset holdings) and (ii) the negative externality of carbon emissions captured by SCC > 0 (trading off clean vs. dirty investments). The climate-hedging premium, reducing expected returns, dominates if $\sigma_{i\eta} > k_i \frac{X}{\gamma}$. The third term, which captures the hedging motive, encompasses both the

standard risk aversion motive (modeled by γ) and the negative externality from carbon (encompassed in SCC). In the absence of non-financial income ($Y_0 = 0$) implying SCC $= 0$ and $X = W$, we arrive at the standard CAPM result $R_i - r = \sigma_{iM}\gamma$. It is also important to highlight that the decomposition is independent of the EIS (ϵ) that only steers the consumption but not the portfolio decision underlying the return decomposition.

3. Evidence on carbon intensity in the US stock market

In this section, we document key stylized facts about the carbon intensity of firms traded in the US stock market, examining patterns both at the firm level and across industries. Our dataset consists of emission information for listed firms in the US (ISIN code starting with *US*). The US stock market universe covers roughly 50% of the global equity market. Despite its dominant role in global equity capitalization, the US accounts for about only 14% of worldwide carbon emissions. A substantial share of global emissions originates from large emerging economies such as China (the single largest emitting country), India, and Russia, whose stock markets remain less accessible, however. High pollution levels in these countries largely reflect their large populations. The US nevertheless ranks among the six countries with the highest per-capita emissions, alongside several Gulf states (see Wang and Feng, 2022).

We use total estimated Scope 1 greenhouse gas (GHG) emissions (metric tons of CO_2 equivalents, tCO_2e) provided by LSEG (formerly known as Refinitiv and ASSET4).[17] In the case in which a company provides information on its Scope 1 carbon emissions, we use that data. For other companies the emissions are estimated by Refinitiv. Note that Scope 1 emissions data only contain direct emissions from the company in question. They do not include indirect emissions from energy consumption (included in Scope 2) or along the entire value chain (included in Scope 3). Scope 2 and 3 level data are not only less available but, more importantly, subject to substantial measurement errors and therefore not very informative. By using estimated Scope 1 data we can create a large sample and also avoid potential double counting of emissions. Further, we note that using Scope 1 data is in line with the commonly implemented "polluter pays" principle in carbon taxation. The key stylized facts described below are robust to the use of the reported emissions in Scope 1 instead of the total emissions, and when extending the analysis to Scope 2 emissions (see Appendix F in the Online Supplementary Material).

The data on emissions is available at an annual frequency with the number of firms reporting their emission levels increasing over time (see Figure 3 below). We only use data for US firms that are active and provide carbon information between 2007 —the

[17] Following the greenhouse gas (GHG) protocol (https://ghgprotocol.org), CO_2 equivalent emissions include carbon dioxide (CO_2), methane (CH_4), nitrous oxide (N_2O), hydrofluorocarbons (HFCS), perfluorinated compound (PFCS), sulfur hexafluoride (SF_6) and nitrogen trifluoride (NF_3).

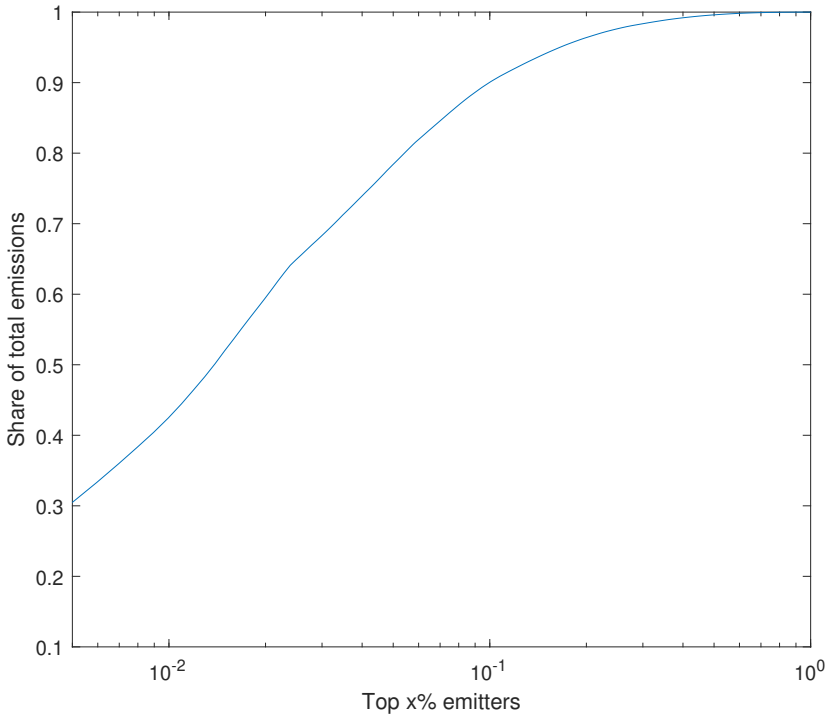


Figure 1. Contribution to overall emissions sorted by top emitter for the year 2019. The top emitters contribute substantially to overall emissions. The 25 largest emitters (roughly 2.5% of all firms) contribute to 2/3 of the carbon emissions. The top 10% share of emitters are responsible for 90% of the emissions. The Gini coefficient – a standard measure of inequality with an upper bound of 1 for positive values as considered here – takes the value of 0.92.

start of the estimated Scope 1 data— and 2019, before the COVID-19 shock that reduced emissions substantially in a transitory manner. All firms in the sample report positive carbon emissions.

Figure 1 displays the contribution to overall emissions sorted by top emitter for the year 2019, the latest year available in our sample. The data suggests that the 25 largest US firm emitters in the sample (roughly 2.5%) contribute to two thirds of the carbon emissions. The top 10% share of emitters is responsible for 90% of the emissions. The Gini coefficient, which is a standard measure of inequality with an upper bound of 1 for positive values as considered here, takes the value of 0.92. Similar results, available upon request, are found for each of the years included in the sample. This is in line with earlier research documenting that roughly two thirds of cumulative carbon emissions are made by only 90 entities world-wide (see Heede, 2014).[18] Overall, the distribution of emitters suggests that by avoiding a few assets, investors can already improve their investment carbon footprint considerably, without substantially reducing the size of the set of firms that they invest in.

[18]Many entities in Heede (2014) are state-owned and therefore not subject to investor influence. Our sample instead focuses on US firms, excluding large state-owned corporations abroad (e.g., UAE, Russia, Iran, Venezuela).

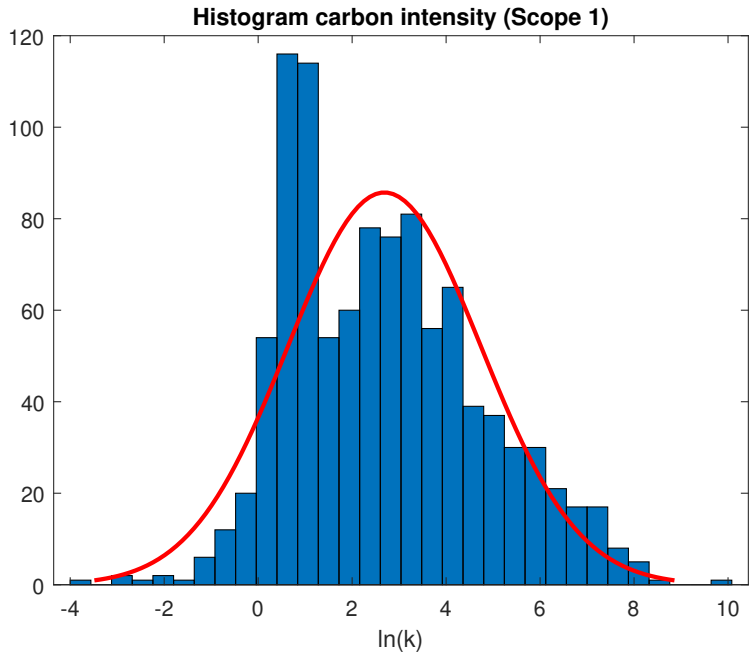


Figure 2. Distribution of carbon intensity k_i (tCO_2e / million USD market value) in 2019. We consider a measure k_i as the carbon footprint metric which maps the amount invested in millions of USD into tCO_2e emitted. There is substantial variation between firms and the measure follows roughly a log-normal distribution ($\log(k) \sim N(\mu, \sigma)$) with $\mu = 2.69$ (0.13) and $\sigma = 2.06$ (0.09) (standard error estimation in brackets).

Next, we construct an estimate of each firm's carbon footprint.[19] In particular, we define the carbon intensity of a firm i at time t as $k_{i,t} = \frac{E_{i,t}}{MV_{i,t}}$ using the emissions $E_{i,t}$ in tCO_2e and the market value $MV_{i,t}$ measured in million USD. The information on market values is obtained from Datastream. Following a pure accounting identity, the market portfolio (index M) is then associated with the total reported emissions ($\sum_i w_{i,t}^M W_t k_{i,t} = \sum_i E_{i,t}$ with $w_{i,t}^M = \frac{MV_{i,t}}{\sum_i MV_{i,t}}$ and $\sum_i MV_{i,t} = W_t$). ca Figure 2 plots the distribution of $\log(k_i)$ for the year 2019. We notice substantial variation of our measure of carbon intensity across firms and conclude that it roughly follows a log-normal distribution with a bit of a skew to the right and with a few large outliers. Similar patterns are found for each of the years in the sample. In Figure 3 we report the mean and median values for k_i over the whole sample period. The plots show that these measures of the carbon intensity vary substantially with a persistent decline over time. The mean carbon intensity is substantially higher than the median which confirms that the distribution is skewed to the right. We also observe that the mean carbon intensity has decreased to roughly half of its value in the beginning of the sample. Nonetheless, it is important to be cautious when interpreting these values for two reasons. First, an increase in climate awareness in financial markets may have led firms to devote more attention to this figure and hence use *creative accounting* to reduce the reported numbers to perform

[19] Throughout, we use "carbon intensity" and "carbon footprint" interchangeably.

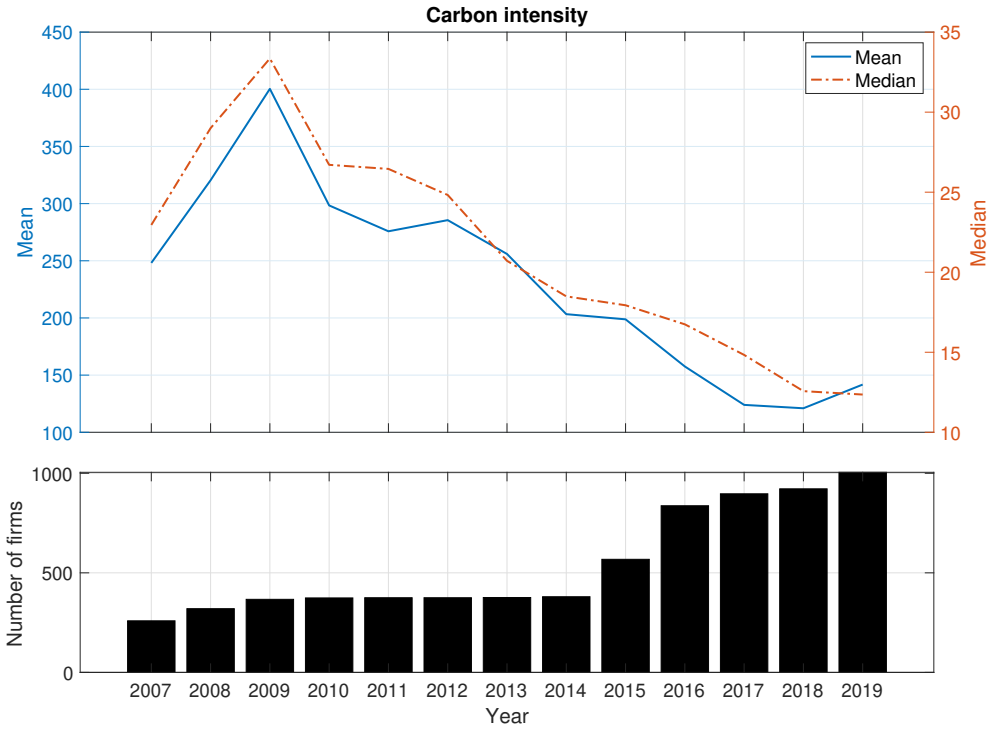


Figure 3. Evolution of carbon intensity k (2007-2019). The top panel plots the time series of firms' carbon intensity, reporting the mean (left axis) and median (right axis). The distribution of carbon intensity is highly right-skewed, with the mean consistently exceeding the median, and it exhibits a gradual decline over time. The bottom panel displays the number of firms with available emissions estimates in each year. Our sample ends in 2019, before the COVID-19 pandemic, which caused a temporary drop in global emissions. The observed downward trend in intensity partly reflects rising market valuations during the period, while reported annual CO_2 emissions roughly halved between 2007 and 2019.

well according to this metric.[20] Second, the stock market valuation increased during the sample period partly explaining a large portion of the observed decrease in this metric.

We also calculate the carbon intensity by industry, grouping firms based on their first two SIC codes leading to nine broad categories. We provide some summary statistics of this metric for the year 2019 in Table 1. In general, industries at the top of the value chain tend to exhibit larger carbon footprints. Based on Scope 1 data, the Transport/Energy sector appears to have a large average carbon intensity, although the dispersion is high within the sector. Further, the Mining and Wholesale Trade sectors are very carbon intense. The high carbon intensity of the latter is, however, likely due to the fact that the sample is skewed upwards by some petrol producers. Interestingly, the financial industry (FIRE) comes out as the cleanest sector, with a very low degree of dispersion. Hence, the current emphasis on green investing promoted by this sector might also direct

[20]For example, in the Online Supplementary Material (Appendix F) we discuss evidence of firms increased Scope 3 emissions which could be a sign that firms outsource their own emissions in this different measure.

Table 1. Descriptive statistics for k by industry for the year 2019. While mining has the highest average k, manufacturing (constituting a large part of the sample) is prone to high dispersion as characterized by the coefficient of variation (CoV). The wholesale trade sector – which among others features petrol producers – is also highly skewed. We do not report the results for the agricultural sector separately since it only features three distinctive companies. FIRE is an acronym for Finance, Insurance and Real Estate sector. Number in brackets report shares of total sample [in %].

Industry	Mean	Std. Dev.	CoV	Min	Max	Observations
Construction	46.23	46.11	1.00	1.27	168.01	16 [2%]
FIRE	2.94	7.30	2.48	0.07	96.32	221 [22%]
Manufacturing	191.78	1,052.96	5.49	0.00	20,212.96	423 [42%]
Mining	499.10	894.20	1.79	19.18	4395.32	25 [2%]
Retail Trade	88.23	255.94	2.90	1.29	2,205.29	78 [8%]
Services	54.22	289.22	5.33	0.10	3,096.02	128 [13%]
Transport/Energy	378.02	626.96	1.66	1.00	3943.17	80 [8%]
Wholesale Trade	354.45	816.21	2.30	0.03	3,746.05	32 [3%]
Total Sample	149.95	752.89	5.02	0.00	20,212.96	1,003 [100%]

the investors' attention towards it. By contrast, the Manufacturing and Services sectors show large within-sector dispersions with notable outliers exhibiting very high carbon intensities. From an investor's perspective, these findings imply not only that it makes sense to look more carefully into the emissions of individual companies, but they also speak against using a standard sector allocation model, for example.

4. Quantitative model outcomes

In this section, we take the Merton–Nordhaus model of Section 2 to the data. We begin by presenting the baseline calibration and use it to estimate the implied social cost of carbon. We then employ the calibrated model to quantify the identified channels and construct optimal portfolios under climate concerns. Finally, we assess how well the optimal portfolios account for the observed returns across US firms in our sample.

4.1 Baseline calibration

Table 2 presents our baseline calibration. The preference parameters are chosen to align with common practices in both the asset pricing literature with recursive preferences and the IAM literature. We set the coefficient of RRA to $\gamma = 10$, following Hambel et al. (2021a), and fix the subjective discount rate at $\rho = 5.8\%$ per year, consistent with the market-based value used by van den Bremer and van der Ploeg (2021) in their baseline calibration. The appropriate value for the EIS is more contentious, reflecting a lack of

Table 2. Baseline calibration.

Parameter	Description	Value
	Preferences	
ρ	Subjective discount rate	5.80%
ϵ	Elasticity of Intertemporal Substitution	0.95
γ	Relative risk aversion	10
	Climate block	
δ	Carbon absorption rate	0.40%
σ_η	Carbon volatility	0.028%
	Economic block	
Y_0	Non-financial income	USD 12.8 trillion
ψ	Damage scaling parameter	5.3×10^{-5}
r	Risk-free rate	0.89%

consensus in the macroeconomics and finance literature.[21] We set the value for the EIS to $\epsilon = 0.95$, a value below one, in line with the arguments put forth by Guvenen (2006, 2009). In robustness checks, we vary these parameters to study the sensitivity of our results (see Appendix E.3 in the Online Supplementary Material for details).

The climate block in the model follows closely the parsimonious two-box carbon model with a permanent and a temporary basin of Golosov et al. (2014) and its extensions (e.g., van der Ploeg and Rezai, 2021).[22] Using the calibrated values in Golosov et al. (2014) we identify an annual carbon absorption rate of $\delta = 0.4\%$.[23] This is also broadly in line with the calibration in van den Bremer and van der Ploeg (2021) who use a modeling approach that is similar to ours (one-box in continuous time) to calibrate $\delta = 0.35\%$ and show that this also fares well compared to more sophisticated carbon models.

The value of ψ, the damage (% income loss) per tonne of stock carbon in the atmosphere, is probably the most controversial in the SCC debate. Following Golosov et al. (2014), we set a baseline value of $\psi = 5.3\%$/Tera tonne stock carbon in the atmosphere in our linear approximation to the damage function. The error between the approximation and the true nonlinear functional form is increasing with the stock of carbon η, with the approximation implying larger damages, thus being more prudent. For a carbon stock in the atmosphere of three Tera tonnes, a value for which all carbon stocks are being completely burnt out, the approximation error is only 1.2% (see Appendix E.2 in the Online Supplementary Material for details).

[21] The EIS remains subject to considerable debate, with estimates ranging from 0.5 to 2.0 (see Cai and Lontzek, 2019 and Jensen and Traeger, 2021).

[22] Recall, however, that η in the model measures CO_2 units. Our calibration instead is for units of carbon and thus the values described reflects the effects of rescaling the model to carbon units by accounting for the atom weight of oxygen and pure carbon.

[23] In their (discrete time) model, a share ϕ_L stays perpetually in the atmosphere, a share $1 - \phi_0$ is immediately removed, and the residual declines geometrically at a rate ϕ: $(1-\delta)^t = \phi_L + (1-\phi_L)\phi_0(1-\phi)^t$.

The model also requires quantifying the size of the stochastic fluctuations in the stock of carbon. For our baseline calibration, we use weekly data on the accumulated levels of atmospheric CO_2 from the Mauna Loa records.[24] We use a rolling-window one-year moving average to filter out the effects of natural seasons. In line with the model, we then take the first difference of the log-transformed data to compute a time series for the relative increments in η. The annual standard deviation of the resulting filtered time series for carbon is $\sigma_\eta = 0.028\%$, a value substantially lower than the one for stock returns shown below.[25] Finally, following the literature on IAMs we set the current carbon stock in the atmosphere, η_0, to a value of roughly 850 GtC (cf. e.g., Cai and Lontzek, 2019).[26]

To calibrate Y_0, the current value of the flow of non-financial income, we set W_0 equal to the value of the combined market capitalization of the US stock and bond markets in 2019. According to the Securities Industry and Financial Markets Association (SIFMA), this amounted to 74.7 trillion USD. We then use a wealth-national income ratio of 3.5 (excluding housing) from Piketty and Zucman (2014) and a capital-income share of 0.4 (Karabarbounis and Neiman, 2013) to compute a value of $Y_0 = 12.8$ trillion USD.[27]

Regarding the investment opportunity set, we set the risk-free rate to $r = 0.88\%$, the average three-month rate from the Fama "Risk-Free Rates" database in CRSP over our sample period. This value is consistent with asset pricing calibrations, typically ranging from $r = 0.6 - 0.8\%$, and somewhat lower in recent studies (Benzoni et al., 2011; Mehra and Prescott, 2003).

For the risky assets, we collect weekly stock prices for the period 2007-2019 of the US companies included in the carbon emission sample of Section 3. The data is obtained from Datastream. To generate sensible results, we restrict our analysis to firms that have been actively traded during the whole sample period and that include information on their carbon emission for at least the last two years. This leaves us with a balanced panel sample of $N = 246$ companies per year. Using this information, we fix the vector of expected returns, $\mathbf{R}$, equal to the average annualized returns over the sample period. Given the large number of assets in the sample, we estimate the covariance matrix of returns, $\mathbf{\Omega}$, using the linear shrinkage estimator of Ledoit and Wolf (2004) to reduce estimation errors and ensure it is positive definitive and thus invertible. In line with the

[24] This is the world largest active volcano situated in Hawaii. The data is available from May 1974 and can be freely accessed at https://gml.noaa.gov/ccgg/trends/data.html.

[25] Using the Mauna Loa records, Hambel et al. (2021a) report an annual value of $\sigma_\eta = 0.78\%$ from a Geometric Brownian Motion process for carbon concentrations. Our estimate is lower, as a result of the filtering process to remove the natural seasonality. Alternatively, van den Bremer and van der Ploeg (2021) report estimates of σ_η that fluctuate between 0.21% and 0.31% for an Arithmetic Brownian Motion specification using the 2,000-year Law Dome Ice Core dataset.

[26] Golosov et al. (2014) set their initial value equal to 802 GtC, an estimate of the atmospheric concentrations in the year 2000.

[27] This figure is calculated as 74.7 trillion USD $\times (1 - 0.4)/3.5 = 12.8$ trillion USD. The data on the combined market capitalization can be retrieved from the SIFMA's 2024 Capital Markets Fact Book, freely accessible at https://www.sifma.org/resources/research/statistics/fact-book/.

Table 3. Annual returns and carbon intensity for assets in numerical experiments. The top panel in the table reports descriptive statistics on the annualized realized equity return and carbon intensity of a sample of $N = 246$ US traded firms over the period 2007-2019. The middle panel reports the average annualized return and carbon intensity sorted by quartiles of carbon intensity, Q_1 to Q_4. The bottom panel reports the average annualized return and carbon intensity across industries in the sample. The (group average) standard deviations for returns and coefficient of variation for carbon intensity are reported in parenthesis below the means in the middle and bottom panels, while the sample sizes (N) within each quartile and industry are in parenthesis, with the respective shares of total sample in brackets. Carbon intensity is measured in tCO_2e per million USD market value.

	Equity returns	**Carbon intensity**		**Carbon/return correlation**
		2007-2019	**2019**	
Panel A: Full sample				
Mean	11.75%	393.63	283.11	2.6%
Std. Dev.	34.03%	1,145.63	970.26	3.0%
Min	-12.80%	0.56	0.05	-4.8%
Max	37.53%	10,721.94	13,047.43	9.7%
Panel B: Across quartiles of carbon intensity				
Q_1 ($N_1 = 61$ [25%])	12.44%	6.53	4.06	2.5%
	(33.81%)	(0.61)	(1.00)	(1.17)
Q_2 ($N_2 = 62$ [25%])	11.62%	27.64	22.36	3.0%
	(30.75%)	(0.30)	(0.74)	(1.10)
Q_3 ($N_3 = 61$ [25%])	11.52%	120.06	131.50	3.1%
	(35.30%)	(0.53)	(1.26)	(0.93)
Q_4 ($N_4 = 62$ [25%])	11.41%	1,409.65	967.55	2.0%
	(36.27%)	(1.39)	(1.82)	(1.42)
Panel C: Across industries				
Transport/Energy ($N = 33$ [13%])	10.54%	1,494.10	730.61	2.0%
	(27.81%)	(1.38)	(1.29)	(1.54)
Mining ($N = 12$ [5%])	8.99%	599.18	305.83	2.1%
	(42.18%)	(1.39)	(1.35)	(1.71)
Manufacturing ($N = 123$ [50%])	12.64%	282.06	280.80	3.0%
	(33.37%)	(3.57)	(4.30)	(0.95)
Wholesale Trade ($N = 6$ [2%])	7.31%	232.51	243.54	1.4%
	(30.56%)	(1.93)	(1.72)	(1.50)
Services ($N = 23$ [9%])	14.02%	92.34	174.56	2.8%
	(38.67%)	(3.62)	(4.19)	(1.14)
Construction ($N = 2$ [1%])	8.84%	88.98	131.73	4.0%
	(36.52%)	(1.28)	(1.40)	(0.26)
Retail trade ($N = 21$ [9%])	11.25%	82.01	71.47	1.9%
	(33.13%)	(1.00)	(1.31)	(1.96)
FIRE ($N = 26$ [11 %])	9.96%	8.73	3.26	2.7%
	(38.53%)	(0.87)	(0.64)	(1.04)

analysis of Section 3, we fix the carbon footprint of the firms in the sample, $\mathbf{k}$, to their value in the year 2019 when computing optimal portfolios. Finally, we also correlate each asset return with the relative change in the accumulation of carbon to obtain an estimate of $\boldsymbol{\rho}_{\eta S}$ (see footnote 6).

Table 3 contains descriptive statistics for these measures. In Panel A, we report

the mean, standard deviation, min and max values for the full sample of firms. We report statistics for carbon intensities for both the complete time sample and the 2019 period only. As already discussed, there is a large dispersion in terms of carbon intensity between firms with an overall decrease in the later years. In contrast, the correlations between asset returns and carbon are small and only fluctuate between -4.8% and 9.7% per year. Together with the covariance matrix of asset returns, they imply $\|\hat{\boldsymbol{\rho}}_{\eta S}\|^2 = 0.54$, suggesting that climate risk is not fully spanned by the financial market.

Panel B, reports the mean return and carbon footprint sorted by quartiles of the carbon intensity. The cross-sectional average of the standard deviation of returns and the coefficient of variation for carbon intensity are reported in parenthesis. As shown earlier, the distribution of carbon intensity is very skewed with extreme observations in the top 25% quartile. We see that firms with lower carbon intensities exhibit slightly higher realized returns as reported in the literature (Bauer et al., 2022; Pástor et al., 2022).

Panel C presents a similar analysis across industries, where the terms in brackets indicate each industry's overall percentage share in the sample. We rank the industries by carbon intensity with Energy/Transport being the dirtiest and the Finance Insurance and Real Estate (FIRE) being the cleanest. Relative to the larger sample of firms in Table 1, the reporting restrictions placed on carbon emissions bias the sample towards heavier emitters with a larger average carbon intensity due to a larger sampling of high-emitting sectors (Transport/Energy, Mining, Manufacturing) and lower sampling of low-emitting sectors (mostly FIRE).

4.2 The social cost of carbon

We begin by considering the case of complete markets and therefore set $\|\hat{\boldsymbol{\rho}}_{\eta S}\|^2 = 1$.[28] Although this assumption implies a correlation structure between climate risk and asset returns that is far from that observed in the data, it enables us to use the closed-form solution of Section 2.4 to examine how the model's predictions respond to variations in key parameter values.

In row (1) of Table 4 we report a SCC of 52.97 USD per tC flow or 14.45 USD per tCO_2e flow, based on our baseline calibration. These values are broadly consistent with SCC estimates reported in the literature (Cai and Lontzek, 2019; van den Bremer and van der Ploeg, 2021; van der Ploeg and Rezai, 2021). Drupp et al. (2024) surveyed scholars conducting work on climate economics on their recommendation on the price of carbon and found a mean SCC of approximately 50 USD per tC, both internationally and for the US. At present, only a few countries —namely, Sweden, Finland, Liechtenstein, and Switzerland— have implemented carbon prices that exceed this level, while most

[28] This is achieved through the transformation $\hat{\boldsymbol{\rho}}_{\eta S}/\|\hat{\boldsymbol{\rho}}_{\eta S}\|^2$, where $\hat{\boldsymbol{\rho}}_{\eta S}$ is computed from the observed correlations in the data.

Table 4. Sensitivity of the social cost of carbon. This table reports the social cost of carbon in USD per tC and USD per tCO_2e for various parameterizations under complete markets ($\|\hat{\boldsymbol{\rho}}_{\eta,S}\|^2 = 1$) and a linear damage function.

	r	δ	ψ	σ_η	**SCC (USD/tC)**	**SCC (USD/tCO_2e)**
(1) Baseline	0.89%	0.40%	5.30%	0.028%	52.97	14.45
(2) Low r	**0.40%**	0.40%	5.30%	0.028%	86.25	23.52
(3) High r	**1.50%**	0.40%	5.30%	0.028%	35.97	9.81
(4) Null δ	0.89%	**0.00%**	5.30%	0.028%	77.04	21.01
(5) Low δ	0.89%	**0.35%**	5.30%	0.028%	55.13	15.04
(6) Low ψ	0.89%	0.40%	**1.10%**	0.028%	10.99	3.00
(7) High ψ	0.89%	0.40%	**20.50%**	0.028%	205.16	55.95
(8) Null σ_η	0.89%	0.40%	5.30%	**0.00%**	52.45	14.30
(9) High σ_η	0.89%	0.40%	5.30%	**0.78%**	73.16	19.95

countries maintain substantially lower prices (Timilsina, 2022).

Rows (2) to (9) of Table 4 illustrate the sensitivity of the SCC to alternative parameterizations of its key determinants. In general, the SCC decreases with the risk-free rate (r) and the carbon absorption (δ), while it increases with the sensitivity of the damage function (ψ) and the volatility of the stock of carbon (σ_η). More specifically, the table shows that using an interest rate equal to 1.50%, corresponding to the discount rate used in Stern (2007), roughly halves the SCC. The trend to lower risk-free rates, say to a level of 0.6% as that used in Benzoni et al. (2011), on the contrary raises the SCC. Similar increases in the SCC can be obtained by lowering the rate of carbon absorption in the atmosphere, which effectively reduces the discount rate. The case of no carbon depreciation in row (4) of Table 4 mirrors the dynamics for the stock of carbon used in Dietz and Venmans (2019). This specification, which is consistent with recent evidence suggesting that emissions lead to an immediate increase global temperatures, implies a higher SCC. The results further highlight the critical role of the damage elasticity ψ. Under the worst-case scenario considered in Golosov et al. (2014), the SCC rises to approximately 205 USD per tC flow. Greater variability in the carbon stock, as in Hambel et al. (2021a), also leads to larger SCC values. When $\sigma_\eta = 0$, as shown in row (8), carbon concentrations evolve deterministically and are therefore risk-free, implying complete markets. In this case, the risk-free SCC, given by Equation (2.33), equals 52.45 USD per tC flow, close to the baseline value due to the low volatility of the observed carbon concentration.

In addition to the above findings, Appendix C of the Online Supplementary Material analyzes a two-asset version of the Merton-Nordhaus model —green (g) and brown (b)— to analytically characterize how the SCC depends on both the correlation between assets, $\mathbf{C} = [\rho_{ij}]_{i,j=g,b}$ (cf. footnote 6), and the correlation between assets and carbon concentrations, $\boldsymbol{\rho}_{\eta S} = [\rho_{\eta g}, \rho_{\eta b}]^\top$. Our analysis shows that when climate damages tend to materialize in "bad times" for investors, the marginal value of mitigation is higher. A

more negative correlation between returns and changes in the carbon stock, particularly for the green asset ($\rho_{\eta g} < 0$), raises the SCC because climate shocks tend to arrive when portfolios are already performing poorly, calling for more aggressive pricing of the externality. This mechanism is evident in counterfactual scenarios with negative return-carbon comovement, which generate higher SCC values than those with positive comovement.

The ability to hedge climate risk across assets plays a central role. When green and brown returns move in opposite directions (i.e., low or negative ρ_{gb}), diversified portfolios cannot offset climate exposures, making the SCC more sensitive to return–carbon correlations. As shown in the illustrations in Appendix C of the Online Supplementary Material, lowering ρ_{gb} increases the responsiveness of the SCC to the sign and magnitude of $\rho_{\eta i}$. By contrast, when ρ_{gb} is high and positive, market risk dominates, cross-asset hedging is limited, and the SCC reacts less to $\rho_{\eta i}$. These patterns are amplified by higher climate volatility σ_η: for larger σ_η, changes in the correlation structure generate materially larger movements in the SCC than under the low-volatility baseline.

Quantitatively, however, these correlation-driven effects are second order for our baseline calibration. Empirically, σ_η is small and the estimated return–carbon correlations $\rho_{\eta i}$ are small and slightly positive. Consequently, the implied hedge premiums are negligible (cf. Table C.1 of the Online Supplementary Material) and the SCC reported in the main text is well approximated by the values in Table 4. The green-brown dichotomy in Appendix C of the Online Supplementary Material therefore primarily serves to bound and explore the sensitivity of the SCC to alternative correlation structures.

4.3 Portfolio composition with climate concerns

In this section, we examine optimal portfolio allocations when climate externalities are explicitly taken into account. We compare alternative portfolio strategies, with and without climate considerations, and assess their performance along several dimensions. The portfolio allocations examined here are intrinsically normative. They describe the choices of a climate-conscious investor given current prices and the SCC implied by the model. They are not equilibrium market holdings, which are determined only after imposing market clearing. Their role is to illustrate how climate externalities influence optimal investment decisions. We highlight two complementary approaches that address portfolio choice under climate concerns.

First, we apply the simple allocation rule derived in Proposition 4 as a simple target function for portfolio construction. This result provides a tractable allocation rule that captures the trade-off between financial return and climate impact in the absence of carbon-concentration volatility. Under Corollary 1, the rule extends naturally to settings with short-sale constraints, requiring non-negative positions in risky assets. This approach offers a practical tool for constructing climate-sensitive portfolios without solving

the full nonlinear optimization problem.

Second, we compute optimal portfolios by solving the Merton–Nordhaus model under incomplete markets and binding short-sale constraints, which corresponds to more realistic market environments in which portfolio choice and climate-hedging motives interact. In this case, we solve the model numerically using a finite-difference approximation of the maximized HJB equation (2.9) on a two-dimensional grid over financial wealth and carbon concentration (see Munk and Sørensen, 2010; Achdou et al., 2022). Further details are provided in Appendix D of the Online Supplementary Material.[29]

For reference, we also report portfolios from the closed-form solution in Section 2.4. Although this benchmark abstracts from several frictions, it provides a point of comparison. As shown in Appendix E.2 of the Online Supplementary Material, the closed-form portfolios remain remarkably robust: optimal allocations differ only marginally when (i) climate risk is unspanned, (ii) damages are nonlinear, or (iii) short sales are restricted.

Table 5 summarizes key financial performance measures and environmental indicators across a variety portfolio strategies. Panel A, reports the portfolio's expected risk premium, volatility, standard Sharpe ratio (SR), and carbon-adjusted Sharpe ratio (SR*). We also provide each portfolio's implied carbon footprint measured in tonnes of CO_2e per million USD of market value (tCO_2e/MUSD). Panels B and C display the (normalized) portfolio composition by carbon intensity quartiles and by industry, respectively.

Column (1) of Table 5 reports the performance of an investment strategy in which each stock is weighted by its share of carbon emissions in the investment opportunity set. The allocation of this carbon-share portfolio spans the entire universe of assets and, as expected, is heavily concentrated in the highest carbon-intensity quartile Q_4. Industry exposures display a similar skew, dominated by Transport/Energy, followed by Manufacturing, while other sectors play only a minor role. This carbon-intensive benchmark, delivers a SR of 0.50 and a SR* of 0.374 due to its large carbon footprint of 1,487.83 tCO_2e/MUSD.

Columns (2) and (3) report the results of a naive diversification rule and a market-based strategy, respectively. Under naive diversification, a fraction $1/N$ of wealth is allocated to each of the N risky assets. The market-based strategy corresponds to a value-weighted portfolio that can be passively implemented via an Exchange Traded Fund. The naive portfolio delivers a SR of 0.552 and SR* of 0.531, with a carbon footprint of 283.11 tCO_2e/MUSD. Its weights are evenly distributed across carbon-intensity quartiles and tilted toward Manufacturing, with moderate exposures to Transport/Energy and Services. By contrast, the market-based portfolio achieves stronger performance (SR = 0.808, SR* = 0.801) and a much lower carbon footprint (82.85 tCO_2e/MUSD), while being concentrated in the lowest carbon-intensity quartiles and maintaining only limited allocations in high-emission sectors such as Transport/Energy and Mining.

[29] Using a Chebyshev projection method (Gaspar and Judd, 1997; Judd, 1992; Parra-Alvarez, 2018) as a robustness check yields very similar allocations. Results available upon request.

Table 5. Financial and environmental characteristics of alternative US portfolio strategies. This table reports the financial and environmental characteristics of various heuristic and optimized portfolio strategies, along with their composition across industries (sorted by decreasing carbon intensity) and carbon-intensity quartiles. Portfolio weights are normalized to 100% ($\widehat{w}_i = \frac{w_i}{\sum_{i=1}^{N} w_i}$) within the universe of $N = 246$ risky assets from Table 3. All figures are annual. "Carbon share" refers to weighting by firms' carbon shares; "MS" denotes the maximum Sharpe ratio; "MSC" the maximum carbon-adjusted Sharpe ratio; and "MN" the optimal strategies from the baseline Merton–Nordhaus model. A dagger (†) indicates strategies with no short-sale constraints, and a double dagger (‡) strategies under spanning and no short-sale constraints (closed-form). "CF" is the portfolio's carbon footprint (tCO_2e per million USD); "Temp." is the implied increase in global temperature over the next 80 years (in °C), with 90% confidence intervals in brackets; "Share assets" is the fraction of the 246 assets with portfolio weights exceeding 10^{-6}; and "SCC" is the Social Cost of Carbon.

	Carbon Share (1)	**Naive** $(1/N)$ (2)	**Market** (3)	**MS**† (4)	**MSC**† (5)	**MN**‡ (6)	**MN**† (7)	**MS** (8)	**MSC** (9)	**MN** (10)	**MN** High SCC (11)	**MN** High Hedge (12)
Panel A: Financial and environmental performance metrics												
$\mathbb{E}[R_p - r]$	8.5%	10.9%	13.6%	94.5%	90.3%	90.3%	90.3%	21.1%	21.2%	21.2%	21.7%	21.3%
σ_p	17.1%	19.7%	16.9%	25.3%	24.7%	24.7%	24.7%	14.6%	14.7%	14.7%	15.1%	14.8%
SR	0.500	0.552	0.808	3.732	3.659	3.659	3.659	1.441	1.441	1.441	1.436	1.441
SR⋆	0.374	0.531	0.801	3.825	3.900	3.900	3.900	1.434	1.435	1.435	1.423	1.435
CF	1,487.83	283.11	82.85	-1,625.89	-4,117.03	-4,117.03	-4,117.19	68.28	52.29	52.28	34.74	51.26
Temp. (°C)	–	3.28 (1.68, 4.80)	1.99 (1.02, 2.92)	–	–	–	–	1.90 (0.97, 2.78)	1.80 (0.92, 2.63)	1.80 (0.92, 2.63)	1.68 (0.86, 2.47)	1.79 (0.92, 2.62)
Share assets	100%	100%	100%	100%	100%	100%	100%	7%	7%	6%	7%	6%
SCC	52.97	52.97	52.97	52.97	52.97	52.97	52.98	52.97	52.97	53.01	205.48	62.39
Panel B: Portfolio composition across quartiles of carbon intensity (%)												
Q_1	1.67	24.80	41.63	-36.31	-34.48	-34.48	-34.48	27.45	28.23	28.23	30.34	28.31
Q_2	9.69	25.20	37.32	177.94	189.63	189.62	189.63	56.55	58.04	58.04	61.44	58.29
Q_3	11.24	24.80	11.85	10.88	49.05	49.06	49.06	0.00	0.00	0.00	0.00	0.00
Q_4	77.40	25.20	9.20	-52.51	-104.20	-104.20	-104.20	16.00	13.73	13.73	8.22	13.41
Panel C: Portfolio composition across industries (%)												
Trans./Energy	50.01	13.41	10.36	0.29	-13.93	-13.93	-13.93	11.23	9.90	9.90	5.61	9.64
Mining	4.21	4.88	1.71	-7.29	0.76	0.76	0.76	0.00	0.00	0.00	0.00	0.00
Manuf.	37.90	50.00	46.96	92.91	95.92	95.92	95.92	53.28	53.72	53.72	56.44	53.82
Whsle. Trade	0.30	2.44	0.67	0.04	5.24	5.24	5.24	0.00	0.00	0.00	0.00	0.00
Services	1.77	9.35	12.97	-49.64	-49.84	-49.84	-49.84	0.00	0.00	0.00	0.00	0.00
Construction	0.06	0.81	0.11	-27.13	-25.63	-25.63	-25.63	0.00	0.00	0.00	0.00	0.00
Retail Trade	5.30	8.54	15.83	95.68	87.96	87.96	87.96	30.41	31.16	31.16	32.41	31.29
FIRE	0.46	10.57	11.39	-4.86	-0.48	-0.48	-0.48	5.08	5.22	5.22	5.53	5.24

We next turn to optimized portfolios, beginning with a strategy that maximizes the SR. Column (4) of Table 5 reports the unconstrained solution. Without short-sale constraints, the portfolio consists of extreme long–short positions across assets, generating a very high SR and a negative carbon footprint, reflecting large short positions in carbon-intensive firms. This should not be mistaken for a carbon sink that removes carbon from the atmosphere. Column (8) presents the corresponding long-only optimal portfolio, which still outperforms the market. It achieves this by overweighting Retail Trade at the end of the value chain and underweighting the FIRE sector. Interestingly, despite not incorporating any climate considerations, the long-only portfolio exhibits a lower carbon footprint (68.28 tCO_2e/MUSD) than the market (82.85 tCO_2e/MUSD). This portfolio is also more concentrated than the market, involving fewer holdings and thus greater exposure to idiosyncratic risk. Overall, the unconstrained and long-only SR optimal portfolios stand in sharp contrast to the previous non-optimized strategies: they achieve markedly higher risk-adjusted returns while also delivering lower carbon footprints. Importantly, none of the portfolios so far incorporate climate externalities; they provide a benchmark for evaluating portfolios that internalize climate externalities.

The remaining columns report optimized strategies that explicitly target the carbon impact. Column (5) presents the unconstrained optimal portfolio obtained by maximizing the $SR^\star$, which penalizes portfolio returns by the carbon footprint weighted with the SCC (set to the baseline value in Table 4). Absent short-sale constraints, the optimal strategy consists of extreme long–short positions, tilting heavily long Retail/Manufacturing and short high-emission sectors. Imposing short-sale constraints in Column (9) leaves financial performance essentially unchanged relative to the constrained SR portfolio in Column (8) but materially lowers the carbon footprint, from 68.28 to 52.29 tCO_2e/MUSD, raising $SR^\star$ slightly. Compositionally, the long-only $SR^\star$ portfolio concentrates around 86% in the two lowest carbon-intensity quartiles, modestly underweights Transport/Energy, and increases Retail and FIRE. Relative to the unconstrained SR portfolio in Column (4), the unconstrained $SR^\star$ solution delivers a similar SR but achieves a far more negative carbon footprint by shorting carbon-intensive firms more aggressively. Across both constrained and unconstrained settings, $SR^\star$-based strategies provide similar financial performance as their SR counterparts, while reducing carbon footprints by roughly 23%.

Finally, we consider the optimal portfolio arising from the Merton-Nordhaus model. Proposition 4 shows that in the unconstrained case, when the hedging demand is negligible (e.g., when $\sigma_\eta = 0$), the optimal portfolio coincides with the one that maximizes the $SR^\star$. This prediction is confirmed by the results: the unconstrained full-model portfolio in Column (7) is nearly identical to the unconstrained $SR^\star$ portfolio in Column (5). For reference, Column (6) reports the closed-form solution from Corollary 2, obtained under full spanning and no short-sale constraints. As expected, this case exhibits the same extreme long–short structure and implausibly high performance measures. Importantly,

the implied SCC is nearly identical across the two approaches.

Corollary 1 extends this result to the case of short-sale constraints. The constrained Merton–Nordhaus solution in Column (10) implies a SCC of 53.01 USD/tC, which is close to the value in the unconstrained model. Thus, the SCC is also relatively robust to investment constraints of this type. The optimal portfolio closely matches the constrained SR* portfolio in Column (9). Both have the same financial performance and a carbon footprint of about 52.3 tCO_2e/MUSD, with almost identical allocations concentrated in the two lowest carbon-intensity quartiles, overweighting Retail and Manufacturing while underweighting Transport/Energy. These close similarities are fully consistent with the theory: given the very small calibrated value of σ_η, hedging motives are negligible, and the carbon-adjusted Sharpe ratio provides an accurate approximation to the fully optimized portfolio both with and without short-sale constraints.

To assess the robustness of our results, Appendix E.1 of the Online Supplementary Material examines how the numerical solution of the Merton-Nordhaus model responds to parameter changes affecting the SCC, mirroring the experiments in Table 4. The SCC displays a sensitivity pattern similar to that of the closed-form case, and portfolio composition and performance remain remarkably stable. Columns (11) and (12) illustrate two representative scenarios. The *High SCC* case ($\psi = 20.50\%$/TtC) yields an extreme SCC value of 205.48 USD/tC shifting weights toward cleaner assets, away from carbon-intensive Transport/Energy firms, and lowering the portfolio's carbon footprint to 34.74 tCO_2e/MUSD. The *High Hedge* case, which raises carbon volatility ($\sigma_\eta = 0.78\%$) and risk aversion ($\gamma = 20$ under CRRA preferences) induces only minor reallocations due to limited hedging opportunities: most assets are weakly but positively correlated with carbon growth (cf. Table 3). Overall, optimal portfolio weights differ by less than 5% across industries and less than 6% across carbon-intensity quartiles relative to the baseline portfolio in Column (10).

As a thought experiment, we also assess the potential impact of the different portfolio strategies on global warming (relative to pre-industrial times) if the investor were also to determine the flow of carbon emissions. We rely on the simple model of Matthews et al. (2009), which assumes a linear relationship between temperature increases and cumulative emissions, $\Delta T_t = \phi \eta_t$.[30] Given the considerable uncertainty surrounding ϕ, we follow Barnett et al. (2020) and report temperature effects using both the central estimate and a range covering the top and bottom 5% of the estimation distribution (in parenthesis). The results are reported in Table 5, where we display the implied increase in global temperatures by the year 2100—roughly the lifespan of a newborn today. Both the market-based and optimized portfolios are associated with expected temperature increases below 2.5°C,

[30] This model abstracts from atmospheric carbon depreciation (δ) and from technological progress that reduces carbon intensity, thereby biasing estimates upward. It also abstracts from stochastic fluctuations in carbon ($\sigma_\eta = 0$).

a threshold often cited as critical for avoiding climate tipping points. By contrast, the naive diversification strategy breaches this threshold. Even the most carbon-conscious portfolio (Column 8) falls short of the 1.5°C target of the Paris Agreement. These estimates should be interpreted with caution: in particular, scaling US portfolios to the global level using financial market size likely understates total world carbon emissions.

Overall, the portfolio experiments reveal three key insights. First, standard mean –variance optimization alone delivers substantial improvements in financial performance relative to heuristic benchmarks. Climate-adjusted strategies yield comparable risk–return outcomes while reducing portfolio carbon footprints by roughly 20–25%, showing that significant climate impact reductions can be obtained without sacrificing financial efficiency. These gains stem from systematic portfolio reallocation toward cleaner assets and sectors.

Second, the carbon-adjusted Sharpe ratio provides a simple yet powerful allocation rule that closely matches the portfolios obtained from the fully optimized structural model, both with and without short-sale constraints. This close alignment reflects the limited role of hedging motives under empirically realistic levels of climate risk and highlights the implementability of the rule in practice in settings where solving the full dynamic optimization problem is infeasible. The micro-founded optimal portfolio of our structural model effectively reduces to maximizing a carbon-adjusted Sharpe ratio—a familiar problem for practitioners and one that can be solved using standard portfolio optimization routines. As with the conventional Sharpe ratio, this approach does not require specifying individual preferences (ρ, γ, ϵ) and can therefore be delegated to investment managers. Investors can then combine these portfolios with the risk-free asset according to their own preferences, while the only exogenous input required to construct the portfolio itself is the SCC.

Third, the closed-form solution offers a strong analytical benchmark for interpreting and validating the numerical results. Appendix E of the Online Supplementary Material complements the main analysis by relaxing the key assumptions behind the closed form (spanning, linear damages, and the absence of short-sale constraints) and by varying preference parameters, including the "ethical" discount rate of 1.5% common in IAMs, different values of the EIS below and above one, and the special case of CRRA preferences, as well as the interaction between the residual carbon risk and the coefficient of relative risk aversion. Across these exercises, the SCC, optimal consumption, and portfolio allocations deviate only marginally from the predictions of the numerical solution. The main exception arises in scenarios with high risk aversion and substantial climate volatility, which increase the SCC and tilt portfolios further toward low-carbon assets. These results underscore the robustness of the closed-form solution and its usefulness as a benchmark for understanding the full model.

Beyond the main set of portfolio experiments, we conduct two additional robustness exercises that complement the analysis. The first focuses on the emissions data used.

The baseline analysis relies on estimated Scope 1 emissions, which provide a direct measure of firms' own carbon output. Appendix F of the Online Supplementary Material replicates the analysis using combined Scope 1 and Scope 2 emissions. The results are virtually identical to those reported in Table 5, confirming that our main conclusions are not sensitive to this data choice. Using Scope 1 and Scope 2 data jointly, however, introduces potential double counting, as emissions from energy use are attributed both to energy producers and to the non-energy firms that consume that energy. This adjustment also reduces the sample size from $N = 264$ to $N = 196$, as the Wholesale Trade and Construction sectors drop out entirely.

The second exercise evaluates portfolio performance out of sample. All results reported in Table 5 are based on in-sample (IS) evaluation. Appendix G of the Online Supplementary Material assesses the out-of-sample (OOS) properties of the portfolios. Consistent with the literature, noisy mean return estimates lead to weaker OOS Sharpe ratios. In contrast, carbon footprint estimates remain stable: portfolios optimized using the carbon-adjusted Sharpe ratio continue to display substantially lower carbon footprints both in and out of sample.

4.4 Return decomposition

Using Proposition 3, we decompose the time-average realized returns reported in Table 3—used as proxies for expected equity returns when forming optimal portfolio decisions absent short-sale constraints— into three distinct components: (i) the market risk premium, corresponding to the standard CAPM factor, (ii) the carbon externality premium, proportional to the firm's carbon intensity and reflecting compensation for the externality, and (iii) the climate-hedge premium, compensating for exposure to innovations in the carbon concentration process.[31] Table 6 presents this decomposition for portfolios sorted by carbon intensity quartiles and by industry, computed under unspanned climate risks and no short-sale restrictions. Premiums are expressed in basis points, and the contribution of each component to the total expected excess return is shown in parentheses below each figure.

For carbon intensity-sorted portfolios, the decomposition follows the pattern implied by Proposition 3: as carbon intensity raises from Q_1 to Q_4, the carbon externality premium increases systematically, while the large market premium declines. Intuitively, as portfolios become more carbon-intensive, a larger fraction of their expected excess return reflects compensation for the externality rather than exposure to the aggregate market factor. The climate-hedge premium is always negative but economically negligible across all quartiles, remaining well below one basis point in absolute value, consistent with the

[31] In the numerical implementation, we follow the proof of Proposition 3 and set the market portfolio equal to the optimal portfolio $\mathbf{w}$. The CAPM covariance is then obtained directly from the return covariance matrix as $\sigma_{iM} = (\mathbf{\Omega w})_i$.

Table 6. Return decomposition. Decomposition of returns following Proposition 3 into (i) the market risk premium (CAPM), (ii) the carbon externality premium, and (iii) the climate-hedging premium. Portfolios are sorted by carbon-intensity quartiles and by industry (ordered from highest to lowest carbon intensity). Returns are reported in basis points, with the relative contribution of each component shown in parenthesis. The "High SCC" calibration sets $\psi = 20.5\%$, while the "High Hedge" calibration sets $\sigma_\eta = 0.78\%$ and $\gamma = 1/\epsilon = 20$ (CRRA preferences).

	Baseline			**High SCC**			**High Hedge**		
	CAPM	Carbon	Hedge	CAPM	Carbon	Hedge	CAPM	Carbon	Hedge
Panel A: Returns across quartiles of carbon intensity									
Q_1	1168	0	-0.007	1166	2	-0.031	1168	1	-0.488
	(100%)	(0%)	(0%)	(100%)	(0%)	(0%)	(100%)	(0%)	(0%)
Q_2	1145	3	-0.007	1137	10	-0.031	1145	3	-0.482
	(100%)	(0%)	(0%)	(99%)	(1%)	(0%)	(100%)	(0%)	(0%)
Q_3	976	13	-0.010	939	49	-0.044	974	15	-0.690
	(99%)	(1%)	(0%)	(95%)	(5%)	(0%)	(99%)	(2%)	(-1%)
Q_4	890	147	-0.006	469	568	-0.025	864	173	-0.383
	(86%)	(14%)	(0%)	(45%)	(55%)	(0%)	(83%)	(17%)	(0%)
Panel B: Returns across industries									
Transport/Energy	859	106	-0.005	556	408	-0.022	841	124	-0.344
	(89%)	(11%)	(0%)	(58%)	(42%)	(0%)	(87%)	(13%)	(0%)
Mining	765	44	-0.008	638	171	-0.036	758	52	-0.554
	(95%)	(5%)	(0%)	(79%)	(21%)	(0%)	(94%)	(6%)	(0%)
Manufacturing	1134	41	-0.009	1017	157	-0.036	1127	48	-0.565
	(97%)	(3%)	(0%)	(87%)	(13%)	(0%)	(96%)	(4%)	(0%)
Wholesale	607	35	-0.004	506	136	-0.016	601	41	-0.256
	(95%)	(5%)	(0%)	(79%)	(21%)	(0%)	(94%)	(6%)	(0%)
Services	1288	25	-0.009	1215	98	-0.040	1284	30	-0.618
	(98%)	(2%)	(0%)	(93%)	(7%)	(0%)	(98%)	(2%)	(0%)
Construction	775	19	-0.012	721	74	-0.053	773	22	-0.817
	(98%)	(2%)	(0%)	(91%)	(9%)	(0%)	(97%)	(3%)	(0%)
Retail	1026	10	-0.004	996	40	-0.019	1024	12	-0.290
	(99%)	(1%)	(0%)	(96%)	(4%)	(0%)	(99%)	(1%)	(0%)
FIRE	906	0	-0.009	905	2	-0.036	907	1	-0.561
	(100%)	(0%)	(0%)	(100%)	(0%)	(0%)	(100%)	(0%)	(0%)

low volatility of carbon concentrations and the small and slightly positive correlations between asset returns and carbon concentrations reported in Table 3. Quantitatively, the estimated carbon premium between the highest (Q_4) and lowest (Q_1) emitters is around 1.47% per year in the baseline scenario. This difference reflects the gap in average carbon intensities of approximately $\Delta k = 1{,}403$ tCO_2/MUSD between the top and bottom quartiles, multiplied by the SCC.

Across industries, a similar pattern emerges. Carbon-intensive sectors such as Transport/Energy exhibit higher carbon externality premiums and slightly lower market risk premiums, whereas low-emission sectors show the opposite. The climate-hedge premium remains uniformly negative and very small, with the Construction sector displaying the largest (most negative) hedge premium due to its relatively strong positive covariance with carbon emissions, therefore generating higher returns in periods of pronounced emissions.

For instance, the gap between Transport/Energy and the low-emission FIRE sector implies a carbon premium of roughly 1.06% per year, consistent with their average carbon intensity difference of about 1,485 tCO_2/MUSD.

We further examine the robustness of these patterns to alternative calibrations. Under the *High SCC* scenario, the carbon externality premium becomes more pronounced, dominating the total expected excess return for Q_4 portfolios (exceeding 50% of the total premium). The climate-hedge premium also increases slightly in absolute terms but remains well below one basis point. Under the *High Hedge* scenario[32], the climate-hedge premium is markedly pronounced but still below one basis point, while the SCC rises slightly to 62.39 USD/tC, modestly amplifying the carbon externality premium relative to the baseline.

Taken together, these results indicate that for the majority of firms, carbon considerations have limited asset pricing implications, with the market risk premium remaining dominant. However, for firms with very high carbon intensity, particularly those in carbon-intensive sectors, the carbon externality premium can be substantial, accounting for a meaningful share of expected excess returns. The climate-hedge premium, by contrast, is uniformly small in economic terms across all calibrations.

5. Conclusion

This paper integrates Merton's Intertemporal Capital Asset Pricing Model with the Integrated Assessment Model logic of Nordhaus to study how climate externalities shape optimal consumption, portfolio choice and equilibrium returns. In our framework, carbon emissions arise both from investors' endogenous portfolio decisions and from exogenous stochastic fluctuations. By linking portfolio decisions to carbon emissions, we identify the Social Cost of Carbon (SCC) as a key determinant of both optimal allocations and equilibrium pricing, effectively internalizing climate damages into financial decision making.

We derive a four-fund separation result showing that optimal portfolios can be expressed as a combination of a risk-free asset, a myopic fund, an externality-pricing fund, and a climate-hedging fund. Building on this characterization, we derive equilibrium returns in the spirit of Merton's ICAPM. In equilibrium, expected excess returns decompose into a market risk premium, a carbon externality premium proportional to the SCC, and a climate-hedge premium. Consistent with empirical evidence, firms that are beneficial (harmful) to the climate earn negative (positive) alphas relative to the standard CAPM.

Using US stock market data for 2007–2019, we estimate an implied SCC of about 53 USD per tC, consistent with benchmark values from the climate-economics literature.

[32]In Appendix E of the Online Supplementary Material, we show that the effect of risk aversion is even more pronounced under incomplete markets, while the EIS has no impact on portfolio choice and thus on the return decomposition.

Accounting for climate externalities has only a modest effect on optimal portfolios: portfolio weights shift slightly across industries and carbon-intensity quartiles, yet the overall risk–return trade-off remains largely intact. Climate considerations primarily influence required returns through the SCC channel, modestly increasing the cost of capital for carbon-intensive firms while preserving the standard mean–variance structure of portfolio choice. Hedging motives play a negligible role, as the volatility of atmospheric carbon concentrations and their correlations with asset returns are both small. Short-sale constraints, by contrast, exert a comparatively larger influence, rendering climate-aware portfolios more concentrated—not necessarily because of stronger environmental motives, but due to limits on hedging brown exposure.

While the quantitative effects are modest, the framework provides a tractable link between damage-based climate economics and portfolio theory, showing how environmental externalities can be incorporated into optimal investment decisions through the SCC. Future work could extend this structure to heterogeneous investors to study the interaction of green and brown investors, heterogeneous beliefs, the exit vs. voice margin, imperfect information about emissions, or policy interventions such as carbon taxes and green subsidies. Empirically, applying the model to more recent data and to markets beyond the US would help assess the robustness and generality of our findings. In this context, our framework offers a theoretical benchmark for evaluating and constructing *green portfolios* that internalize the SCC within a coherent investment and asset-pricing framework.

References

Acemoglu, D., P. Aghion, L. Bursztyn, and D. Hemous (2012): "The Environment and Directed Technical Change," *American Economic Review*, 102, 131–166.

Achdou, Y., J. Han, J.-M. Lasry, P.-L. Lions, and B. Moll (2022): "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach," *The Review of Economic Studies*, 89, 45–86.

Aswani, J., A. Raghunandan, and S. Rajgopal (2023): "Are Carbon Emissions Associated with Stock Returns?" *Review of Finance*, 28, 75–106.

Avramov, D., S. Cheng, A. Lioui, and A. Tarelli (2022): "Sustainable Investing with ESG Rating Uncertainty," *Journal of Financial Economics*, 145, 642–664.

Baker, S. D., B. Hollifield, and E. Osambela (2022): "Asset Prices and Portfolios with Externalities," *Review of Finance*, 26, 1433–1468.

Bansal, R., M. Ochoa, and D. Kiku (2016): "Climate Change and Growth Risks," NBER Working Paper No. 23009, NBER.

Bansal, R. and A. Yaron (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59, 1481–1509.

Barnett, M. (2023): "Climate Change and Uncertainty: An Asset Pricing Perspective," *Management Science*, 69, 7562–7584.

Barnett, M., W. Brock, and L. P. Hansen (2020): "Pricing Uncertainty Induced by Climate Change," *Review of Financial Studies*, 33, 1024–1066.

Barrage, L. (2020): "Optimal Dynamic Carbon Taxes in a Climate-Economy Model with Distortionary Fiscal Policy," *The Review of Economic Studies*, 87, 1–39.

Bauer, M. D., D. Huber, G. D. Rudebusch, and O. Wilms (2022): "Where Is the Carbon Premium? Global Performance of Green and Brown Stocks," *Journal of Climate Finance*, 1, 100006.

Benzoni, L., P. Collin-Dufresne, and R. S. Goldstein (2011): "Explaining Asset Pricing Puzzles Associated with the 1987 Market Crash," *Journal of Financial Economics*, 101, 552–573.

Bolton, P. and M. Kacperczyk (2021): "Do Investors Care about Carbon Risk?" *Journal of Financial Economics*, 142, 517–549.

——— (2023): "Global Pricing of Carbon-Transition Risk," *Journal of Finance*, 78, 3677–3754.

Broccardo, E., O. D. Hart, and L. Zingales (2022): "Exit versus Voice," *Journal of Political Economy*, 130, 3101–3145.

Cai, Y. and T. S. Lontzek (2019): "The Social Cost of Carbon with Economic and Climate Risks," *Journal of Political Economy*, 127, 2684–2734.

Chen, G., B. Mojon, and D. Vayanos (2025): "The Impact of Green Investors on Stock Prices," *Review of Finance*, forthcoming.

Chittaro, L., M. Piazzesi, M. Sena, and M. Schneider (2025): "Asset Returns as Carbon Taxes," NBER Working papers 34342.

Cochrane, J. H. (2021): "Portfolios for Long-Term Investors," *Review of Finance*, 26, 1–42.

Cox, J. C., J. E. Ingersoll, and S. A. Ross (1985a): "A Theory of the Term Structure of Interest Rates," *Econometrica*, 53, 385–407.

——— (1985b): "An Intertemporal General Equilibrium Model of Asset Prices," *Econometrica*, 53, 363–384.

De Angelis, T., P. Tankov, and O. D. Zerbib (2023): "Climate Impact Investing," *Management Science*, 69, 7669–7692.

Dietz, S., C. Gollier, and L. Kessler (2018): "The Climate Beta," *Journal of Environmental Economics and Management*, 87, 258–274.

Dietz, S. and F. Venmans (2019): "Cumulative Carbon Emissions and Economic Policy: In Search of General Principles," *Journal of Environmental Economics and Management*, 96, 108–129.

Drupp, M. A., F. Nesje, and R. C. Schmidt (2024): "Pricing Carbon: Evidence from Expert Recommendations," *American Economic Journal: Economic Policy*, 16, 68–99.

Duffie, D. and L. G. Epstein (1992a): "Asset Pricing with Stochastic Differential Utility," *Review of Financial Studies*, 5, 411–436.

——— (1992b): "Stochastic Differential Utility," *Econometrica*, 60, 353–394.

Epstein, L. and S. Zin (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–69.

Gaspar, J. and K. Judd (1997): "Solving Large-Scale Rational-Expectations Models," *Macroeconomic Dynamics*, 1, 45–75.

Giglio, S., M. Maggiori, K. Rao, J. Stroebel, and A. Weber (2021): "Climate Change and Long-Run Discount Rates: Evidence from Real Estate," *The Review of Financial Studies*, 34, 3527–3571.

Gollier, C. and S. Pouget (2022): "Investment Strategies and Corporate Behaviour with Socially Responsible Investors: A Theory of Active Ownership," *Economica*, 89, 997–1023.

Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski (2014): "Optimal Taxes on Fossil Fuel in General Equilibrium," *Econometrica*, 82, 41–88.

Guvenen, F. (2006): "Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective," *Journal of Monetary Economics*, 53, 1451–1472.

——— (2009): "A Parsimonious Macroeconomic Model for Asset Pricing," *Econometrica*, 77, 1711–1750.

Hambel, C., H. Kraft, and E. Schwartz (2021a): "Optimal Carbon Abatement in a Stochastic Equilibrium Model with Climate Change," *European Economic Review*, 132, 103642.

——— (2021b): "The Social Cost of Carbon in a Non-Cooperative World," *Journal of International Economics*, 131, 103490.

Hambel, C., H. Kraft, and F. van der Ploeg (2024): "Asset Diversifitcation versus Climate Action," *International Economic Review*, 65, 1323–1355.

Hambel, C. and F. van der Ploeg (2025): "Policy Transition Risk, Carbon Premiums, and Asset Prices," *Journal of Monetary Economics*, 152.

Hassler, J., P. Krusell, and A. Smith (2016): *Environmental Macroeconomics*, Elsevier, vol. 2 of *Handbook of Macroeconomics*, chap. 0, 1893–2008.

Heede, R. (2014): "Tracing Anthropogenic Carbon Dioxide and Methane Emissions to Fossil Fuel and Cement Producers, 1854—2010," *Climatic Change*, 122, 229–241.

Heinkel, R., A. Kraus, and J. Zechner (2001): "The Effect of Green Investment on Corporate Behavior," *Journal of Financial and Quantitative Analysis*, 36, 431–449.

Hong, H., N. Wang, and J. Yang (2023): "Welfare Consequences of Sustainable Finance," *The Review of Financial Studies*, 36, 4864–4918.

Hsu, P.-H., K. Li, and C.-Y. Tsou (2023): "The Pollution Premium," *Journal of Finance*, 78, 1343–1392.

Huang, K., H. Zhao, J. Huang, J. Wang, and C. Findlay (2020): "The Impact of Climate Change on the Labor Allocation: Empirical Evidence from China," *Journal of Environmental Economics and Management*, 104, 102376.

Jensen, S. and C. P. Traeger (2021): "Pricing Climate Risk," CESifo Working Paper No. 9196, CESifo.

Judd, K. L. (1992): "Projection Methods for Solving Aggregate Growth Models," *Journal of Economic Theory*, 58, 410–452.

Karabarbounis, L. and B. Neiman (2013): "The Global Decline of the Labor Share," *The Quarterly Journal of Economics*, 129, 61–103.

Karydas, C. and A. Xepapadeas (2022): "Climate Change Financial Risks: Implications for Asset Pricing and Interest Rates," *Journal of Financial Stability*, 63, 101061.

Kreps, D. and E. L. Porteus (1978): "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica*, 46, 185–200.

Ledoit, O. and M. Wolf (2004): "A Well-conditioned Estimator for Large-dimensional Covariance Matrices," *Journal of Multivariate Analysis*, 88, 365–411.

Matthews, H. D., N. P. Gillett, P. A. Stott, and K. Zickfeld (2009): "The Proportionality of Global Warming to Cumulative Carbon Emissions," *Nature*, 459, 829–832.

Mehra, R. and E. C. Prescott (2003): "The Equity Premium in Retrospect," in *Handbook of the Economics of Finance*, ed. by G. Constantinides, M. Harris, and R. M. Stulz, Elsevier, vol. 1 of *Handbook of the Economics of Finance*, chap. 14, 889–938.

Merton, R. C. (1969): "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case," *The Review of Economics and Statistics*, 51, 247–57.

——— (1971): "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," *Journal of Economic Theory*, 3, 373–413.

——— (1973): "An Inter-Temporal Capital Asset Pricing Model," *Econometrica*, 41, 867–887.

Munk, C. and C. Sørensen (2010): "Dynamic Asset Allocation with Stochastic Income and Interest Rates," *Journal of Financial Economics*, 96, 433–462.

Nordhaus, W. D. (1993): "Rolling the 'DICE': An Optimal Transition Path for Controlling Greenhouse Gases," *Resource and Energy Economics*, 15, 27–50.

——— (2015): "Climate Clubs: Overcoming Free-Riding in International Climate Policy," *American Economic Review*, 105, 1339–1370.

Parra-Alvarez, J. C. (2018): "A Comparison of Numerical Methods for the Solution of Continuous-Time DSGE Models," *Macroeconomic Dynamics*, 22, 1555–1583.

Pastor, L., R. F. Stambaugh, and L. A. Taylor (2021): "Sustainable Investing in Equilibrium," *Journal of Financial Economics*, 142, 550–571.

——— (2022): "Dissecting Green Returns," *Journal of Financial Economics*, 146, 403–424.

Pedersen, L. H. (2025): "Carbon Pricing versus Green Finance," *Journal of Finance (forthcoming)*.

Pedersen, L. H., S. Fitzgibbons, and L. Pomorski (2021): "Responsible Investing: The ESG-Efficient Frontier," *Journal of Financial Economics*, 142, 572–597.

Piketty, T. and G. Zucman (2014): "Capital is Back: Wealth-Income Ratios in Rich Countries 1700—2010," *The Quarterly Journal of Economics*, 129, 1255–1310.

Pástor, L., R. F. Stambaugh, and L. A. Taylor (2022): "Dissecting Green Returns," *Journal of Financial Economics*, 146, 403–424.

Sauzet, M. and O. D. Zerbib (2022): "When Green Investors are Green Consumers," Working Paper, Boston University.

Stern, N. (2007): *The Economics of Climate Change: The Stern Review*, Cambridge University Press.

Timilsina, G. R. (2022): "Carbon Taxes," *Journal of Economic Literature*, 60, 1456–1502.

van den Bremer, T., v. d. P. Frederick, and S. Wills (2016): "The Elephant in the Ground: Managing Oil and Sovereign Wealth," *European Economic Review*, 82, 113–131.

van den Bremer, T. S. and F. van der Ploeg (2021): "The Risk-Adjusted Carbon Price," *American Economic Review*, 111, 2782–2810.

van der Ploeg, F. and A. Rezai (2021): "Optimal Carbon Pricing in General Equilibrium: Temperature Caps and Stranded Assets in an Extended Annual DSGE Model," *Journal of Environmental Economics and Management*, 110.

Wang, M. and C. Feng (2022): "Tracking the Inequalities of Global Per Capita Carbon Emissions from Perspectives of Technological and Economic Gaps," *Journal of Environmental Management*, 315, 115144.

Weil, P. (1990): "Nonexpected Utility in Macroeconomics," *The Quarterly Journal of Economics*, 105, 29–42.

Zhang, S. (2025): "Carbon Returns across the Globe," *Journal of Finance*, 80, 615–645.

Appendix

A. Proof of Proposition 3

This proof follows Merton (1973) closely. However, the notation is somewhat different and we work with matrices and vectors as we solve for the risk premiums. Merton (1973) works with scalars and, in the process, he imposes some additional constraints. As such, we provide a slight generalization. Since the time index is not of particular interest, we suppress it throughout.

The value of the market is given by

$$M = \sum_{i=0}^{N} \overline{N}_i S_i, \tag{A.1}$$

where $\overline{N}_i$ is the number of supplied shares in asset i. The zeroth asset is a risk-free asset and assets 1 through N are risky.

The supply of shares is allowed to change over time. Applying Itô's lemma on the above equation, we get

$$\mathrm{d}M = \sum_{i=0}^{N} \overline{N}_i \mathrm{d}S_i + \sum_{i=0}^{N} (S_i + \mathrm{d}S_i)\mathrm{d}\overline{N}_i. \tag{A.2}$$

Here, the first term represents capital gains, while the second term represents an *expansion of the number of shares*.

Since the representative investor holds all the resources, the change in the value of the market portfolio must be equal to the change in the representative investor's wealth, i.e., $dM = dW$, where

$$\mathrm{d}W = W \sum_{i=0}^{N} w_i \frac{\mathrm{d}S_i}{S_i} + (Y - C)\mathrm{d}t. \tag{A.3}$$

According to the analysis in Merton (1973),

$$(Y - C)\mathrm{d}t = \sum_{i=0}^{N} (S_i + \mathrm{d}S_i)\mathrm{d}\overline{N}_i, \tag{A.4}$$

and

$$W \sum_{i=0}^{N} w_i \frac{\mathrm{d}S_i}{S_i} = \sum_{i=0}^{N} \overline{N}_i \mathrm{d}S_i. \tag{A.5}$$

Thus, we have that $w_i = \overline{N}_i S_i / W$. Further, since the representative investor holds all the resources, her wealth is equal to the value of the market portfolio, implying that $w_i = \overline{N}_i S_i / W = \overline{N}_i S_i / M$. The interpretation of Equation (A.4) above is that the

increase in the total market value induced by an expansion in the number of shares exactly matches the savings rate of the representative investor.

As in Merton (1973), we define the price of the market portfolio (S_M) and the number of shares in the market portfolio ($\overline{N}$) through $M = \overline{N}S_M$. Thus, by Itô's lemma,

$$\mathrm{d}M = \overline{N}\mathrm{d}S_M + \mathrm{d}\overline{N}(S_M + \mathrm{d}S_M), \tag{A.6}$$

where

$$\overline{N}\mathrm{d}S_M = \sum_{i=0}^{N} \overline{N}_i \mathrm{d}S_i, \tag{A.7}$$

and

$$\mathrm{d}\overline{N}(S_M + \mathrm{d}S_M) = \sum_{i=0}^{N} \mathrm{d}\overline{N}_i (S_i + \mathrm{d}S_i). \tag{A.8}$$

Equation (A.7) can be restated as

$$\frac{\mathrm{d}S_M}{S_M} = \sum_{i=0}^{n} w_{Mi} \frac{\mathrm{d}S_i}{S_i}, \tag{A.9}$$

where $w_{Mi} = \overline{N}_i S_i / M$ is the market share of asset i.

Letting $\mathbf{w}_M = [w_{M1}, w_{M2}, ..., w_{MN}]^\top$ be the vector of portfolio weights in the market portfolio allocated to the *risky* assets, we have that

$$\frac{\mathrm{d}S_M}{S_M} = (r + \mathbf{w}_M^\top(\mathbf{R} - r\mathbf{1}))\mathrm{d}t + \mathbf{w}_M^\top \mathbf{\Omega}^{1/2} \mathrm{d}\mathbf{Z}_t, \tag{A.10}$$

and we can define quantities like the instantaneously expected rate of return of the market portfolio ($R_M := r + \mathbf{w}_M^\top(\mathbf{R} - r\mathbf{1})$) as well as the instantaneous variance of returns ($\sigma_M^2 := \mathbf{w}_M^\top \mathbf{\Omega^{1/2}} (\mathbf{\Omega^{1/2}})^\top \mathbf{w}_M = \mathbf{w}_M^\top \mathbf{\Omega} \mathbf{w}_M$) and instantaneous covariances with η and stock i, respectively.

From Equation (A.5), it follows that the optimal portfolio vector in Proposition 1 must be equal to the vector of weights of risky assets in the market portfolio ($\mathbf{w}_M$), implying that

$$\mathbf{w}_M = \mathbf{\Omega}^{-1}(\mathbf{R} - r\mathbf{1})\theta + \mathbf{\Omega}^{-1}\mathbf{k}\kappa + \mathbf{\Omega}^{-1}\mathbf{\Omega}^{1/2}\sigma_\eta \eta \hat{\boldsymbol{\rho}}_{\eta S} H. \tag{A.11}$$

Premultiplying the left-hand and the right-hand sides of the above equation by $\mathbf{\Omega}/\theta$ and solving for $\mathbf{R} - r\mathbf{1}$, we obtain

$$\mathbf{R} - r\mathbf{1} = \frac{1}{\theta}\mathbf{\Omega}\mathbf{w}_M - \frac{\kappa}{\theta}\mathbf{k} - \frac{H}{\theta}\mathbf{\Omega}^{1/2}\sigma_\eta \eta \hat{\boldsymbol{\rho}}_{\eta S}. \tag{A.12}$$

From the above, we can establish a system of four equations:

$$R_i - r = \begin{bmatrix} \sigma_{iM} & k_i & \sigma_{i\eta} \end{bmatrix} \begin{bmatrix} \frac{1}{\theta} \\ -\frac{\kappa}{\theta} \\ -\frac{H}{\theta} \end{bmatrix}, \tag{A.13}$$

$$R_M - r = \begin{bmatrix} \sigma_M^2 & k_M & \sigma_{M\eta} \end{bmatrix} \begin{bmatrix} \frac{1}{\theta} \\ -\frac{\kappa}{\theta} \\ -\frac{H}{\theta} \end{bmatrix}, \tag{A.14}$$

$$R_{N-1} - r = \begin{bmatrix} \sigma_{(N-1)M} & k_{N-1} & \sigma_{(N-1)\eta} \end{bmatrix} \begin{bmatrix} \frac{1}{\theta} \\ -\frac{\kappa}{\theta} \\ -\frac{H}{\theta} \end{bmatrix}, \tag{A.15}$$

$$R_N - r = \begin{bmatrix} \sigma_{NM} & k_N & \sigma_{N\eta} \end{bmatrix} \begin{bmatrix} \frac{1}{\theta} \\ -\frac{\kappa}{\theta} \\ -\frac{H}{\theta} \end{bmatrix}, \tag{A.16}$$

where, provided that we define $\mathbf{o}_i$ as an $N \times 1$ vector with a 1 in its ith entry and zeroes everywhere else, $\sigma_{iM} = \mathbf{o}_i^\top \boldsymbol{\Omega} \mathbf{w}_M$, $\sigma_{i\eta} = \mathbf{o}_i^\top \boldsymbol{\Omega}^{\mathbf{1/2}} \sigma_\eta \eta \hat{\boldsymbol{\rho}}_{\eta S}$, $\sigma_M^2 = \mathbf{w}_M^\top \boldsymbol{\Omega} \mathbf{w}_M$, $\sigma_{M\eta} = \mathbf{w}_M^\top \boldsymbol{\Omega} \sigma_\eta \eta \hat{\boldsymbol{\rho}}_{\eta S}$, and $k_M = \mathbf{w}_M^\top \mathbf{k}$.[33] Further, Equation (A.14) follows from premultiplying both sides of Equation (A.12) by $\mathbf{w}_M^\top$.

With $\boldsymbol{\Xi}$ defined as in Proposition 3, the last three equations can be written as

$$\begin{bmatrix} R_M - r \\ R_{N-1} - r \\ R_N - r \end{bmatrix} = \boldsymbol{\Xi} \begin{bmatrix} \frac{1}{\theta} \\ -\frac{\kappa}{\theta} \\ -\frac{H}{\theta} \end{bmatrix}. \tag{A.17}$$

The result follows from pre-multiplying the above by $\boldsymbol{\Xi}^{-1}$, solving for $\begin{bmatrix} \frac{1}{\theta} & -\frac{\kappa}{\theta} & -\frac{H}{\theta} \end{bmatrix}^\top$ and inserting it into Equation (A.13). ■

[33] Note that some quantities result from letting $i = N - 1$ and $i = N$, respectively.