What Is the Carbon Premium a Premium On?*

Marcin Kacperczyk
Imperial College London

April 27, 2026

Abstract

I develop an asset-pricing framework in which climate risk premia depend on the price of risk and how firm characteristics map into cash-flow exposure. Which characteristics are priced is a question about the underlying economic mechanism, not a statistical choice among proxies. A leading example is emission intensity, which is sufficient only under restrictive homogeneity and proportionality conditions. I derive two testable implications, spanning and scale invariance, and reject both in international data: emission intensity fails as a sufficient statistic, with emission levels carrying independent cross-sectional information conditional on firm value. These results point to scale-dependent exposure functions. More broadly, commonly used measures recover projections of underlying exposure, offering a unified explanation for conflicting evidence on climate risk premia and implications for capital allocation.

*I thank Franklin Allen, Harjoat Bhamra, Patrick Bolton, Paolo Fulghieri, Lorenzo Garlappi, Deeksha Gupta, Max Konradt, Magda Rola-Janicka, Zach Sautner, Maxime Sauzet, Luca Taschini, Michela Verardo, Alex Wagner, Jiang Wang, David Zerbib, and seminar participants at CFRN Conference, City University London, FEB-RN, NBS-PRI-ECGI, and UNC-RFS Climate and Energy Conference for helpful comments and discussions. Marcin Kacperczyk is with Imperial College London, CEPR, and ECGI (m.kacperczyk@imperial.ac.uk).

1 Introduction

Investors, regulators, and policy makers increasingly speak about a carbon premium or greenium in asset prices. Getting its measurement right matters: portfolio screens, regulatory taxonomies, and corporate discount rates all depend on correctly identifying which firm attributes carry climate risk. Yet there is little agreement on which characteristics capture priced climate exposure. Empirical studies proxy climate exposure using absolute emissions or changes in emissions (e.g., Bolton and Kacperczyk (2021, 2023)), emissions scaled by firm value or revenue (Pedersen (2025)), sectoral classifications (Pastor et al. (2022)), or composite transition-risk scores (e.g., Sautner et al. (2023), Cenedese et al. (2024), Eskildsen et al. (2024)). Consistent with this lack of agreement, pairwise correlations across commonly used climate risk measures are far from perfect, suggesting that these proxies capture distinct dimensions of underlying exposure rather than a single priced object. The coexistence of these approaches therefore leaves a basic issue unresolved: which firm characteristics actually govern climate risk premia in equilibrium?

This disagreement reflects a deeper model-selection problem. The finance literature has largely evaluated climate characteristics by asking which proxies predict returns. But expected returns are determined by how climate and policy shocks enter firm cash flows. Different economic mechanisms imply different cash-flow exposure functions and therefore different candidate sufficient statistics for pricing climate risk. Absent a mapping from economic mechanisms to normalized exposure, empirical comparisons across alternative measures remain difficult to interpret, as different characteristics recover projections of distinct underlying exposures rather than a common priced object.

I address this model-selection problem by developing a framework that characterizes which firm characteristics can serve as sufficient statistics for pricing climate risk. Climate-related shocks affect firms through dollar cash-flow exposures summarized by exposure functions $g(Z_{i,t})$ that map firm characteristics into dividend sensitivities. Because returns are defined per unit of invested wealth, equilibrium risk premia load on normalized exposures $g(Z_{i,t})/V_{i,t}$, where $V_{i,t}$ denotes firm value. Identifying priced characteristics therefore requires identifying the underlying exposure function implied by the economic mechanism. This perspective resolves the apparent conflict across empirical studies: different emission-based measures recover different projections of $g(Z_{i,t})/V_{i,t}$ onto firm attributes, and which projection is informative depends on which economic channels are active in the sample being studied. Studies documenting premia on different characteristics are therefore not competing proxies for a common priced object but measurements of distinct underlying exposures.

To characterize how market structure determines the equilibrium price of climate risk, I develop a tractable model with heterogeneous investors who face unspanned climate damages and differ in their concern about those damages. The model delivers a simple benchmark in which firms' cash-flow sensitivities to climate or policy shocks scale with emission levels, generating a factor structure in which expected returns reflect value-normalized sensitivities together with an equilibrium price of risk. In this benchmark, the price of risk admits a decomposition into a component that compensates investors for bearing unhedgeable climate risk and a component driven by heterogeneous climate concern.

I apply this framework to emission intensity, the leading policy-relevant and widely used measure of climate risk, to illustrate its implications. This application yields a characterization that emission intensity $X_{i,t}/V_{i,t}$ is a sufficient statistic for pricing climate risk only under two joint restrictions. First, primitive cash-flow sensitivities must scale proportionally with firm size, so that the underlying exposure is homogeneous of degree one in emissions and firm value. Second, priced sensitivities must be invariant across firms with the same emission intensity, implying that two firms with identical intensity carry the same normalized exposure regardless of their scale or other characteristics. Under these joint conditions, expected returns depend only on emission intensity and not separately on the underlying primitives from which it is constructed. These restrictions generate falsifiable cross-sectional implications—spanning and scale invariance—that can be tested directly in the data. More generally, this characterization illustrates how candidate characteristics correspond to specific assumptions about the underlying exposure function.

The framework has two central implications. First, market structure affects the equilibrium *price* of climate risk but not the mapping from firm fundamentals into priced exposure: the extent to which climate risks are spanned and the degree of heterogeneity in climate concern determine the price of risk, while the function $g(Z_{i,t})$ is governed by how shocks enter firms' cash flows. Emission intensity therefore appears as a priced characteristic only under the restrictive conditions that make normalized sensitivities collapse to $X_{i,t}/V_{i,t}$.

Second, misidentifying the relevant exposure function distorts both pricing and real outcomes. When investors or regulators rely on characteristics that do not correspond to $g(Z_{i,t})/V_{i,t}$, equilibrium prices reflect projections onto the chosen metric rather than underlying exposure, directing capital toward firms that appear low-risk under the selected characteristic and inducing firms to optimize measured metrics rather than the underlying drivers of climate risk. Efficiency improvements that raise output and emissions are a concrete illustration: they can reduce reported intensity

while increasing equilibrium exposure to priced shocks.[1] Quantifying these distortions requires first identifying the correct exposure function, precisely the object this paper characterizes.

Building on this structure, I embed the model in a unified linear-factor framework for climate risk that nests transition, physical, and technology channels as distinct factors with separate exposure functions. In this representation, expected excess returns take the form

$$E_t[R_{i,t+1} - R_{f,t}] \;=\; \sum_k \lambda_{k,t} \, \frac{g_k(Z_{i,t})}{V_{i,t}},$$

where each factor k captures a distinct channel of climate risk, such as policy shocks, physical damages, or technology changes, and $g_k(Z_{i,t})$ summarizes how that channel enters firm i's dollar cash flows. This representation provides a systematic way to assess when a given scalar characteristic can summarize the multidimensional exposure $g_k(Z_{i,t})/V_{i,t}$ across factors. In general, many economically natural mechanisms violate the homogeneity and scale-invariance conditions required for intensity sufficiency, so that emission intensity need not summarize priced exposure in equilibrium.

I test these restrictions in the context of emission intensity using an international panel of firms with Scope 1 emissions. Consistent with the model's predictions, both restrictions implied by intensity sufficiency are rejected: the pricing of emission intensity varies systematically with firm size, violating scale invariance, and emission levels carry independent cross-sectional information about expected returns once firm value is controlled for, violating the spanning restriction in pooled tests. Because intensity sufficiency implies that log emissions and log firm value enter returns only through their difference, I test this as a linear coefficient restriction and reject it at horizons of one and twelve months ($p = 0.008$ and $p = 0.031$, respectively). In economic terms, a one-standard-deviation increase in log Scope 1 emissions predicts a 0.36% increase in next-month expected returns (about 4.3% annualized), with an effect of 0.15% per month (1.8% annualized) at the twelve-month horizon. The scale-invariance rejection is mirrored by the cross-sectional pattern of the intensity premium, which is largest among the smallest firms and attenuates monotonically with firm scale. This pattern persists within size quintiles: emission levels retain positive and statistically significant explanatory power in every size bin, and the spanning restriction is rejected within the three smallest quintiles, indicating that the pooled rejection is not an artifact of residual size heterogeneity.

The objective of my analysis is not to identify a best proxy for climate risk premia but to provide testable restrictions that any candidate characteristic must satisfy to serve as a sufficient statistic

[1] This mechanism is related to the Jevons paradox (Jevons, 1865), whereby improvements in efficiency can increase total resource use through behavioral and equilibrium responses.

for priced exposure. The rejection of emission intensity therefore does not imply that another proxy is universally sufficient; it highlights the need to discipline measurement by the economic mechanisms that determine exposure.

A natural implication of this framework is that aggregating across multiple measures can obscure the underlying economic object being priced. Such aggregation reduces measurement error only if the underlying signals correspond to the same economic object. When different measures instead originate from distinct economic mechanisms, such as policy exposure, technology choice, regulatory compliance, or investor preferences, aggregation obscures interpretation rather than improving it. In that case, it becomes unclear what economic exposure the resulting score is intended to proxy.

Finally, the framework clarifies the role of emission intensity in policy discussions. In the sustainable-discount-rate literature, intensity emerges naturally as a sufficient statistic for implementing policy objectives because cash-flow distortions scale proportionally with emissions. This rationalizes intensity's prevalence in regulation and policy design but clarifies its scope: it emerges as the appropriate measure when a planner applies a uniform per-unit shadow price on emissions across firms–not as a general feature of asset pricing equilibria.

Overall, the paper makes three contributions. First, it reframes the measurement of climate risk premia as a problem of identifying normalized exposure functions $g(Z_{i,t})/V_{i,t}$ that map firm characteristics into priced cash-flow sensitivities. This perspective implies that commonly used characteristics are informative only to the extent that they recover the underlying exposure function, rather than serving as primitives of the pricing relation. Second, it provides a general characterization of when a given scalar characteristic is a sufficient statistic for priced exposure. In particular, a characteristic must recover $g(Z_{i,t})/V_{i,t}$ up to scale, which imposes sharp and testable cross-sectional restrictions: spanning and scale invariance. These restrictions translate the distinction between characteristics and exposure functions into empirically implementable tests. Third, it shows that economically natural policy and technology mechanisms generate exposure functions that violate these restrictions, so that commonly used measures need not summarize priced risk. As a result, empirical proxies for climate risk reflect projections of underlying exposure onto selected characteristics rather than the exposure itself. This misalignment introduces systematic measurement error and can distort capital allocation, including in settings where investors or regulators rely on composite ESG measures that aggregate across heterogeneous exposure channels.

Related Literature. This paper relates to several strands of the climate-finance and sustainable-investing literature. Theoretically, it relates to equilibrium models with heterogeneous ESG preferences. Pastor et al. (2021) develop a static equilibrium in which an exogenous greenness characteristic enters investors' utility and generates a green factor and greenium in expected returns. Related models with ESG-motivated investors show how such preferences affect equilibrium portfolios and expected returns (Pedersen et al., 2021; Zerbib, 2022; De Angelis et al., 2023; Dangl et al., 2025). These papers take the relevant ESG or environmental metric as a primitive of investors' utility or screening rule. My approach is logically prior: regardless of whether premia arise from risk compensation, hedging demand, or preference channels, the cross-sectional mapping from firm characteristics into priced exposure is governed by the cash-flow exposure function $g(Z_{i,t})/V_{i,t}$, which is determined by technology and policy rather than investor preferences. The conditions under which a chosen metric recovers this object are the focus of this paper.

A related literature studies sustainable-discount-rate and social-cost-of-carbon frameworks. Pedersen (2025) shows that in a complete-markets planner economy, the wedge implementing the social optimum is proportional to firms' emission intensity. Chittaro et al. (2025) show that in a frictionless representative-agent setting, a carbon tax is equivalent to a return schedule increasing linearly in emission intensity. Empirically, Pastor et al. (2025a) introduce a firm's carbon burden–the present value of social costs of future emissions–and document a positive association with the cost of capital. In my framework, the carbon burden is a planner-implementation object: SCC-weighted emissions emerge as a sufficient statistic precisely when the homogeneity and scale-invariance conditions of Proposition 3 hold by construction of the implementation wedge (Section 7.3). Whether such objects survive as priced characteristics in equilibrium asset pricing is a distinct question, and the tests in Section 6.2 reject the corresponding restrictions in international data.

A growing literature emphasizes firm-level exposure to climate shocks through cash flows and production decisions (Bolton and Kacperczyk, 2023; Hsu et al., 2023; Cenedese et al., 2024), documenting that carbon-related characteristics affect returns and firm valuations through transition-risk channels. This work establishes that firm-level exposure is empirically relevant; my contribution is to characterize when observable firm characteristics provide sufficient statistics for that exposure, and to derive testable restrictions that discipline the mapping from characteristics to priced risk.

Empirically, a large literature documents correlations between carbon characteristics and expected returns, using emission levels and changes thereof (Bolton and Kacperczyk, 2021, 2023), emission intensity (Pastor et al., 2022; Eskildsen et al., 2024), composite ESG scores, and text-based

climate exposure measures (Sautner et al., 2023). In my framework, each of these corresponds to a projection of a potentially multidimensional exposure onto a particular attribute subset. Pastor et al. (2025b) survey the sustainable-investing literature and identify as an open question which ESG attributes are actually priced in equilibrium. I address this question by characterizing when a given attribute recovers the underlying exposure function and by providing empirical restrictions that such characteristics must satisfy. A related challenge is that ostensibly similar ESG characteristics can be constructed very differently across data providers (Berg et al., 2022). My framework provides a structural interpretation of this divergence: when multiple economic mechanisms jointly determine cash-flow exposure, no scalar index summarizes priced exposure in general, and different providers' weighting choices correspond to projections of the same multidimensional object onto different attribute subsets (Section 7.2). Eskildsen et al. (2024) construct a composite green score by averaging across emission-based measures to mitigate measurement error; my framework implies that such aggregation is informative only when the underlying signals correspond to a common priced object, a restriction that the spanning evidence in Section 6.2 does not support.

More broadly, earlier work on the pricing of "sin" characteristics (Hong and Kacperczyk, 2009) examined settings in which the relevant attribute was determined exogenously by industry classification, allowing the analysis to focus on the mechanism generating premia. The climate setting differs because the appropriate characteristic is itself an open question: emissions, intensity, changes, and trajectory-based measures are all defensible candidates corresponding to distinct underlying mechanisms. I take a model-disciplined approach to selection among them: which emission-based characteristics can legitimately serve as sufficient statistics for exposure to priced climate factors? The framework delivers testable restrictions—scale invariance and spanning—that discipline when emission intensity suffices for equilibrium pricing.

The remainder of the paper is organized as follows. Section 2 develops a tractable incomplete-markets environment with heterogeneous investors, specifying firms, cash flows, and preferences in the presence of unspanned climate risk. Sections 3 and 4 build on this environment to characterize equilibrium pricing in a general linear-factor framework and derive the conditions under which a given characteristic, such as emission intensity, is sufficient to summarize priced exposure. Section 5 studies economically motivated specifications that violate these conditions, illustrating how different mechanisms generate distinct priced characteristics. Section 6 tests the model's cross-sectional restrictions–spanning and scale invariance–using international firm-level data. Section 7 discusses implications for measurement, capital allocation, and policy. Section 8 concludes.

2 A Simple Model with Heterogeneous Investors and Climate Risk

Consider a two-period economy, with $t = 0, 1$. Uncertainty is represented by a probability space Ω with probability measure $\mathbb{P}$. The economy features a single physical consumption good that serves as numeraire, N firms indexed by $i = 1, \ldots, N$, and H households indexed by $h = 1, \ldots, H$. In addition to the N risky assets, a risk-free asset in zero net supply with gross return $R_f > 0$ is available.

2.1 Emissions, climate shock, and damages

Each firm i has a fixed level of emissions $X_i \geq 0$ at date 0, which determines its exposure to the climate shock realized at date 1. Aggregate emissions are

$$\bar{X} \equiv \sum_{i=1}^{N} X_i. \tag{2.1}$$

The economy is exposed to a one-dimensional aggregate climate or policy shock $\zeta : \Omega \to \mathbb{R}$ with

$$\mathbb{E}[\zeta] = 0, \qquad \text{Var}(\zeta) = \sigma_\zeta^2 > 0. \tag{2.2}$$

I assume ζ is normally distributed and all payoffs are jointly normal.

Climate damages, Γ, are linear in ζ and proportional to aggregate emissions:

$$\Gamma(\zeta) \equiv \kappa \bar{X}\, \zeta, \qquad \kappa > 0. \tag{2.3}$$

Equation (2.3) is a reduced-form representation of the idea that the realization of the climate shock ζ scales damages in proportion to $\bar{X}$. The linear specification serves as a convenient benchmark, which I relax in Internet Appendix A by allowing for nonlinear and more general damage specifications.

2.2 Firms and cash flows

Firm i issues a single share of equity in fixed supply. Its date-1 dividend in state $\omega \in \Omega$ is

$$D_i(\omega) = \mu_i + \eta X_i \zeta(\omega), \tag{2.4}$$

where $\mu_i \in \mathbb{R}$ is a deterministic component and $\eta > 0$ is a technology parameter common across firms. The shock ζ is best interpreted as aggregate transition-risk news–unexpected changes in carbon pricing, regulation, enforcement intensity, or clean-technology adoption–under which the linear loading on emissions admits a direct reading: higher-emission firms experience larger cash-flow impacts from per-unit policy shocks, and a more permissive regulatory state simultaneously raises aggregate damages and brown cash flows.[2] Physical climate risk is accommodated only as a reduced-form approximation in this benchmark, because physical damages typically affect firms through heterogeneous channels such as aggregate productivity, demand, supply-chain disruption, and policy responses that need not scale with emissions and need not co-move positively with brown cash flows. The multi-factor extension in Internet Appendix A treats physical, policy, and technology shocks as distinct factors with separate exposure functions and is the appropriate vehicle for analyzing physical risk in its own right. Equation (2.4) should therefore be viewed as a transition-risk benchmark that isolates a simple exposure channel in which cash-flow sensitivity is proportional to emissions. Later sections relax this assumption and allow the exposure function $g(Z_{i,t})$ to depend on richer firm characteristics, illustrating how alternative economic mechanisms generate different priced characteristics.

I further denote V_i as the date-0 market value of firm i's equity, and define the gross return

$$R_i(\omega) \equiv \frac{D_i(\omega)}{V_i}. \tag{2.5}$$

2.3 Households, preferences, and market structure

Household h has date-0 financial wealth $W_0^h > 0$. At $t = 0$, the household chooses a portfolio $\theta^h = (\theta_1^h, \ldots, \theta_N^h)^\top$ in the N risky assets and invests an amount b^h in the risk-free asset, subject to the budget constraint

$$\sum_{i=1}^{N} V_i \theta_i^h + b^h = W_0^h. \tag{2.6}$$

[2]The scalar ζ is signed so that $\zeta > 0$ simultaneously raises aggregate damages through $\Gamma(\zeta) = \kappa \bar{X} \zeta$ and high-emission firms' cash flows through $\eta X_i \zeta$, consistent with an interpretation of ζ being weak-policy news. A more accommodating regulatory environment both allows damages to accumulate and benefits carbon-intensive firms. Under this convention, cash flows of brown firms covary positively with non-tradable damages, which delivers the negative preference component $\lambda_{X,t}^H$ in Section 3.3 and its interpretation as a hedging channel. The alternative convention with ζ being policy tightening, adopted for the carbon-price example in Section 4.1, where $D_{i,t+1} = \tilde{D}_{i,t+1} - \tau_{t+1} X_{i,t}$, is isomorphic after flipping the sign of η and yields identical equilibrium pricing relations.

Date-1 consumption of household h in state ω is then given by

$$C_1^h(\omega) = R_f\, b^h + \sum_{i=1}^{N} \theta_i^h D_i(\omega), \tag{2.7}$$

where R_f denotes the gross risk-free return.

Households have exponential (CARA) preferences over climate-adjusted consumption,

$$U^h = \mathbb{E}\left[-\exp\left(-\alpha_h\big(C_1^h(\omega) - \phi_h \Gamma(\zeta)\big)\right)\right], \tag{2.8}$$

where $\alpha_h > 0$ is the coefficient of absolute risk aversion and $\phi_h \geq 0$ indexes household h's climate concern. The term $\phi_h \Gamma(\zeta)$ captures disutility from aggregate climate damages $\Gamma(\zeta)$, with $\phi_h = 0$ corresponding to a purely financial (brown) investor and $\phi_h > 0$ to a climate-conscious (green) investor. $\Gamma(\zeta)$ should be interpreted as the *non-pecuniary* component of climate damages that affect household utility without flowing through the date-1 budget constraint. Examples of such welfare-relevant externalities include biodiversity loss, health impacts, intergenerational burdens, or sea-level rise. This is distinct from aggregate consumption risk: damages that materially reduce C_1 would enter through dividends and the resource constraint rather than through the preference term. Modeling climate damages as non-pecuniary is what generates the incomplete-markets structure below, since households cannot construct portfolios of traded claims that replicate disutility from externalities they do not bear through their own consumption. The heterogeneity in ϕ_h then governs the extent to which these unpriced welfare effects are internalized in equilibrium portfolio choice.

A canonical example of unspanned climate risk arises from global damages that affect households' welfare but are not directly traded in financial markets. Sea-level rise, biodiversity loss, or health damages from extreme weather may reduce welfare without mapping one-to-one into the cash flows of traded firms. Investors therefore care about these outcomes but cannot construct portfolios that perfectly hedge them. This wedge between welfare-relevant climate damages and traded payoffs generates an incomplete-markets pricing kernel.

Because these damages are non-tradable and investors differ in (α_h, ϕ_h), markets are incomplete with respect to climate risk and the stochastic discount factor that prices traded assets reflects a risk tolerance-weighted aggregation of heterogeneous marginal rates of substitution rather than the planner's sustainable discount factor.[3] The resulting priced climate characteristic is therefore an

[3] I use the term incomplete markets as shorthand for environments in which climate damages are non-tradable and heterogeneously valued across investors. Formally, the climate externality reflects a missing market, while heterogene-

equilibrium object, a function of firm exposures, and the associated premium admits a decomposition into a risk-based brown component and a preference-driven green component, as formalized in Section 3.1.

3 Equilibrium Pricing with Unspanned Climate Risk

This section develops the core framework of the paper. Climate and transition risk affect asset prices through their impact on firms' cash flows, and expected returns are determined by value-normalized exposure functions that map firm characteristics into sensitivities to climate and policy shocks. I formalize this mapping using exposure functions $g(Z_{i,t})$ that translate firm characteristics into dollar sensitivities to climate and policy shocks. This representation accommodates multiple, overlapping channels through which climate risk operates. Emissions, technology, demand conditions, and policy exposure can all affect how firms load on aggregate shocks, and different economic mechanisms imply different exposure functions. As a result, no single characteristic, such as emission intensity, is guaranteed to summarize priced exposure a priori. The framework yields a linear factor representation in which expected returns depend on normalized exposure functions. This representation provides a systematic way to assess when a given characteristic is sufficient for pricing climate risk and when it is not.

3.1 Market incompleteness and non-tradable climate damages

The only traded risky payoffs are the N stocks with dividends given by (2.4). There is a single aggregate climate shock ζ, and I allow for additional idiosyncratic shocks in μ_i, but for the derivations below I focus on the one-factor structure (2.4).

Building on the environment in Section 2, markets are incomplete in that there is no portfolio of traded assets whose payoff replicates the damage term $\Gamma(\zeta)$. Equivalently, ζ is not fully spanned by the space of traded payoffs once the non-pecuniary climate term in preferences is taken into account. Put differently, markets are complete for financial payoffs, but the climate damage term in utility is non-tradable, so the SDF is not the MRS of someone who internalizes the externality in a planner sense.

Nevertheless, equilibrium asset pricing remains well defined. There exists a unique stochastic discount factor m that prices all traded assets, but the incomplete spanning of $\Gamma(\zeta)$ implies that

ity in concern implies that the resulting damage risk is unspanned by traded assets, leading to an incomplete-markets pricing kernel.

this pricing kernel cannot be interpreted as the marginal rate of substitution of a representative planner who internalizes the externality. This distinction from the complete-markets benchmark is central for understanding how climate risk is priced in equilibrium.

3.2 Equilibrium pricing of climate risk

Mean-variance representation Under CARA-exponential utility and joint normality, maximizing expected utility (2.8) is equivalent to maximizing a mean-variance objective. Substituting (2.3), (2.4), and (2.7), we have

$$C_1^h - \phi_h \Gamma(\zeta) = b^h R_f + \sum_{i=1}^{N} \theta_i^h \mu_i + \Big(\sum_{i=1}^{N} \theta_i^h \eta X_i - \phi_h \kappa \bar{X} \Big) \zeta.$$

Thus

$$C_1^h - \phi_h \Gamma(\zeta) = A_h + B_h \zeta,$$

where

$$A_h \equiv b^h R_f + \sum_{i=1}^{N} \theta_i^h \mu_i, \qquad B_h \equiv \sum_{i=1}^{N} \theta_i^h \eta X_i - \phi_h \kappa \bar{X}.$$

The mean-variance representation of (2.8) becomes

$$U^h = -\exp\Big(-\alpha_h \Big(\mathbb{E}[A_h] - \tfrac{\alpha_h}{2} \mathrm{Var}(B_h \zeta) \Big) \Big),$$

so household h chooses (θ^h, b^h) to maximize

$$\mathbb{E}[C_1^h - \phi_h \Gamma(\zeta)] - \frac{\alpha_h}{2} \mathrm{Var}[C_1^h - \phi_h \Gamma(\zeta)]. \tag{3.1}$$

Since b^h only enters linearly and serves to satisfy the budget constraint, the key choice is the vector θ^h.

Using $\mathbb{E}[\zeta] = 0$ and $\mathrm{Var}(\zeta) = \sigma_\zeta^2$, objective (3.1) becomes

$$\sum_{i=1}^{N} \theta_i^h (\mu_i - V_i R_f) \; - \; \frac{\alpha_h}{2} \sigma_\zeta^2 \Big(\sum_{i=1}^{N} \theta_i^h \eta X_i - \phi_h \kappa \bar{X} \Big)^2.$$

First-order conditions and equilibrium prices The first-order condition for θ_i^h is

$$\mu_i - V_i R_f - \alpha_h \sigma_\zeta^2 \Big(\sum_{j=1}^{N} \theta_j^h \eta X_j - \phi_h \kappa \bar{X} \Big) \eta X_i = 0, \qquad i = 1, \ldots, N. \tag{3.2}$$

Define

$$\delta_h \equiv \sum_{j=1}^{N} \theta_j^h \eta X_j - \phi_h \kappa \bar{X}.$$

Then (3.2) is equivalently

$$\mu_i - V_i R_f = \alpha_h \sigma_\zeta^2 \delta_h \eta X_i, \qquad i = 1, \ldots, N. \tag{3.3}$$

Since the left-hand side does not depend on h, equilibrium requires $\alpha_h \sigma_\zeta^2 \delta_h$ to be common across households with interior holdings.

Proposition 1. *In any competitive equilibrium with $\eta > 0$, there exists a scalar λ such that*

$$\mu_i - V_i R_f = \lambda \eta X_i, \qquad i = 1, \ldots, N, \tag{3.4}$$

and for every household h,

$$\delta_h = \frac{\lambda}{\alpha_h \sigma_\zeta^2}. \tag{3.5}$$

Proof is in Appendix A.1.

Equation (3.4) is a pricing relation in dollar units. Dividing by V_i yields the corresponding return representation,

$$\mathbb{E}[R_i] - R_f = \lambda \eta \frac{X_i}{V_i}, \tag{3.6}$$

where $R_i \equiv D_i / V_i$ and $\mathbb{E}[D_i] = \mu_i$. Emission levels X_i govern the primitive cash-flow exposure in dividends, while expected returns load on the value-normalized exposure X_i / V_i. The coefficient λ is the equilibrium price of climate-related risk.

Solving for the price of climate risk Aggregating the definition of δ_h across households, multiplying by η, and imposing market clearing $\sum_{h=1}^{H} \theta_j^h = 1$ yields

$$\eta^2 \bar{X} \; - \; \kappa \eta \bar{X} \sum_{h=1}^{H} \phi_h \; = \; \eta \sum_{h=1}^{H} \delta_h. \tag{3.7}$$

Using (3.5), the right-hand side is $\eta \sum_{h=1}^{H} \delta_h = (\eta\lambda/\sigma_\zeta^2) \sum_{h=1}^{H}(1/\alpha_h)$. Solving for λ gives:

Proposition 2. *In the model above, the equilibrium price of climate risk is*

$$\lambda = \frac{\sigma_\zeta^2 \bar{X}}{\sum_{h=1}^{H} \frac{1}{\alpha_h}} \Big(\eta - \kappa \sum_{h=1}^{H} \phi_h\Big). \tag{3.8}$$

Proof is in Appendix A.1.

Under CARA preferences, λ depends on the distribution of risk aversion (α_h) and aggregate climate concern $\sum_h \phi_h$, and is independent of initial wealth (W_0^h).[4] Combining (3.8) with (3.6) shows how equilibrium expected returns reflect both the normalized exposure X_i/V_i and the aggregate price of climate risk λ, which depends on the distribution of risk aversion and climate concern across investors.

3.3 Decomposing transition risk and climate-concern hedging effects

The equilibrium price of climate risk admits a decomposition that separates compensation for bearing transition risk from a hedging component driven by heterogeneous concern about non-tradable climate damages. This decomposition isolates how market structure affects the price of risk, while the underlying exposure mapping remains governed by how shocks enter firms' cash flows.

Using (3.8), the price of climate risk in the incomplete-markets economy can be written as

$$\lambda_{X,t} = \frac{\sigma_\zeta^2 \bar{X}_t \eta}{\sum_{h=1}^{H} \frac{1}{\alpha_h}} - \frac{\sigma_\zeta^2 \bar{X}_t \kappa}{\sum_{h=1}^{H} \frac{1}{\alpha_h}} \sum_{h=1}^{H} \phi_h.$$

Define

$$\lambda_{X,t}^{\mathrm{TR}} \equiv \frac{\sigma_\zeta^2 \bar{X}_t \eta}{\sum_{h=1}^{H} \frac{1}{\alpha_h}}, \qquad \lambda_{X,t}^{\mathrm{H}} \equiv -\frac{\sigma_\zeta^2 \bar{X}_t \kappa}{\sum_{h=1}^{H} \frac{1}{\alpha_h}} \sum_{h=1}^{H} \phi_h, \tag{3.9}$$

so that

$$\lambda_{X,t} = \lambda_{X,t}^{\mathrm{TR}} + \lambda_{X,t}^{\mathrm{H}}. \tag{3.10}$$

In this representation, $\lambda_{X,t}$ denotes the equilibrium price per unit of emissions exposure $X_{i,t}$. Note that because dividend exposure to the climate shock is proportional to $\eta X_{i,t}$, the technology parameter η is absorbed into the transition-risk component $\lambda_{X,t}^{\mathrm{TR}}$ through the cash-flow loading,

[4] In Internet Appendix B, I show that with CRRA utility the analogous expression features wealth weights given by local risk tolerances.

while the hedging component $\lambda_{X,t}^{\mathrm{H}}$ reflects preference-driven demand for assets that covary with non-tradable climate damages and does not depend on firm-level technology.

The term $\lambda_{X,t}^{\mathrm{TR}}$ is the price of climate risk that would prevail if all households were climate-neutral, that is, if $\phi_h = 0$ for all h. It compensates investors for bearing unhedgeable exposure to the aggregate climate or policy shock ζ_{t+1} through the technological channel $D_{i,t+1} = \mu_{i,t+1} + \eta X_{i,t}\zeta_{t+1}$.

By contrast, $\lambda_{X,t}^{\mathrm{H}}$ arises from heterogeneous concern about non-tradable climate damages and is strictly decreasing in aggregate concern $\sum_h \phi_h$ when $\kappa > 0$. Differentiating (3.9) yields

$$\frac{\partial \lambda_{X,t}}{\partial \phi_h} = -\frac{\sigma_\zeta^2 \bar{X}_t \kappa}{\sum_{j=1}^{H} \frac{1}{\alpha_j}} < 0.$$

This comparative static reflects that assets with greater exposure to the climate shock provide a hedge against non-tradable climate damages, so stronger aggregate climate concern reduces the required compensation for holding such assets.

Given the decomposition (3.10), the expected excess return on firm i can be written as

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \lambda_{X,t}^{\mathrm{TR}} \frac{X_{i,t}}{V_{i,t}} + \lambda_{X,t}^{\mathrm{H}} \frac{X_{i,t}}{V_{i,t}}. \tag{3.11}$$

The first term reflects compensation for bearing non-diversifiable transition (or physical) climate risk, while the second term captures the pricing effect of hedging motives arising from heterogeneous climate concern.

An equivalent decomposition applies to the stochastic discount factor. Writing

$$m_{t+1} = \exp\big(-A_t - B_t \zeta_{t+1}\big), \qquad B_t = \frac{\lambda_{X,t}}{\sigma_\zeta^2},$$

we obtain

$$B_t = B_t^{\mathrm{TR}} + B_t^{\mathrm{H}},$$

with

$$B_t^{\mathrm{TR}} = \frac{\bar{X}_t \eta}{\sum_{h=1}^{H} \frac{1}{\alpha_h}}, \qquad B_t^{\mathrm{H}} = -\frac{\bar{X}_t \kappa}{\sum_{h=1}^{H} \frac{1}{\alpha_h}} \sum_{h=1}^{H} \phi_h.$$

The first component describes how the pricing kernel responds to the climate shock in the absence of climate concern, while the second captures the additional tilt induced by hedging demand.

Overall, this decomposition shows that transition risk and climate-concern hedging effects are conceptually distinct but jointly determine the observed carbon premium. The preference compo-

nent reflects hedging motives arising from non-tradable climate damages rather than direct preferences for holding green assets, although both channels may coexist in practice.

3.4 When is emission intensity a unique sufficient statistic?

The previous subsection shows that expected excess returns are governed by firms' factor exposures, which may depend on a rich set of characteristics $Z_{i,t}$. There is therefore no general reason why any single primitive, such as emission levels $X_{i,t}$ or emission intensity $X_{i,t}/V_{i,t}$, must be sufficient for climate-related pricing. This subsection formalizes and characterizes the restrictive conditions under which emission intensity is a unique sufficient statistic for the priced exposure to a given climate or transition factor.

Definition 3.1. *Consider date t and a one-factor environment in which expected excess returns satisfy*

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \Lambda_t \chi_{i,t},$$

for some scalar Λ_t and some scalar characteristic $\chi_{i,t}$ that may depend on firm attributes. We say that emission intensity *is a sufficient statistic at date t if for any two firms i and j,*

$$\frac{X_{i,t}}{V_{i,t}} = \frac{X_{j,t}}{V_{j,t}} \quad \Rightarrow \quad \chi_{i,t} = \chi_{j,t}.$$

We say that intensity is a unique *sufficient statistic if, in addition, any other characteristic $\tilde{\chi}_{i,t}$ that generates an equivalent one-factor representation must be proportional to $X_{i,t}/V_{i,t}$ in the cross-section, that is, there exists a scalar ψ_t such that*

$$\tilde{\chi}_{i,t} = \psi_t \frac{X_{i,t}}{V_{i,t}} \quad \textit{for all } i.$$

Equivalently, sufficiency requires that $g(Z_{i,t})/V_{i,t}$ depend on firm characteristics only through $X_{i,t}/V_{i,t}$. The formulation using $\chi_{i,t}$ is useful because the empirical question is which observable proxies recover this object.

One-factor environment and value normalization. For clarity, let us suppress time subscripts when no confusion arises. Let f_{t+1} denote a scalar transition or climate factor with

$\mathbb{E}_t[f_{t+1}] = 0$, and consider dividends of the form

$$D_{i,t+1} = \mu_{i,t+1} + g(Z_{i,t}) f_{t+1}, \tag{3.12}$$

where $Z_{i,t}$ collects firm characteristics (including in particular $(X_{i,t}, V_{i,t})$), and $g(\cdot)$ is a deterministic function. The object $g(Z_{i,t})$ is a *cash-flow* exposure (in dividend units). Let the stochastic discount factor admit the linear projection representation

$$m_{t+1} = a_t + b_t f_{t+1} + \varepsilon_{t+1}, \tag{3.13}$$

with $a_t > 0$, b_t $\mathcal{F}_t$–measurable, and ε_{t+1} conditionally mean-zero and orthogonal to f_{t+1}. Using the Euler equation $\mathbb{E}_t[m_{t+1} R_{i,t+1}] = 1$ and the orthogonality of ε_{t+1}, expected excess returns satisfy

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \lambda_t \frac{g(Z_{i,t})}{V_{i,t}}, \qquad \lambda_t = -\frac{b_t \mathrm{Var}_t(f_{t+1})}{\mathbb{E}_t[m_{t+1}]}, \tag{3.14}$$

which follows directly from the covariance between firm returns and the pricing kernel. Market structure determines b_t and hence λ_t, but not the functional form of $g(\cdot)$.[5]

Knife-edge conditions for intensity sufficiency. The next result provides a necessary and sufficient characterization of when the return-relevant exposure $g(Z_{i,t})/V_{i,t}$ depends on firms only through emission intensity $X_{i,t}/V_{i,t}$, and when this dependence is unique up to scale. The result is model-agnostic and does not rely on the incomplete-markets structure developed earlier; it follows from the mapping between dollar cash-flow exposures and returns.

Proposition 3. *Consider the one-factor environment* (3.12)-(3.14) *at a given date* t*. Suppose that:*

(i) The cross-section of firms is rich *in the sense that the set*

$$\mathcal{S}_t \equiv \left\{ \left(\frac{X_{i,t}}{V_{i,t}}, V_{i,t} \right) : i = 1, \ldots, N \right\}$$

exhibits independent variation in emission intensity and firm scale.[6]

[5] In this framework, normalized exposure $g(Z_{i,t})/V_{i,t}$ corresponds to the firm's beta with respect to the underlying shock.

[6] Formally, let $\mathcal{F}_t$ denote the limiting empirical distribution of $\left(\frac{X_{i,t}}{V_{i,t}}, V_{i,t} \right)$ as $N \to \infty$. Condition (i) requires that the support of $\mathcal{F}_t$ contains an open subset of $\mathbb{R}^2_+$, ruling out degenerate cases in which the cross-section of firms concentrates on a lower-dimensional manifold.

(ii) *The exposure function $g(\cdot)$ is homogeneous of degree one in $(X_{i,t}, V_{i,t})$ in the sense that, for all $\lambda > 0$ and all firms i,*

$$g(Z_{i,t}; \lambda X_{i,t}, \lambda V_{i,t}) = \lambda\, g(Z_{i,t}; X_{i,t}, V_{i,t}).$$

(iii) *For any two firms i and j with the same intensity, $X_{i,t}/V_{i,t} = X_{j,t}/V_{j,t}$, their normalized exposures coincide,*

$$\frac{g(Z_{i,t})}{V_{i,t}} = \frac{g(Z_{j,t})}{V_{j,t}}.$$

Then there exists a (measurable) function $\varphi_t : \mathbb{R}_+ \to \mathbb{R}$ such that

$$\frac{g(Z_{i,t})}{V_{i,t}} = \varphi_t\left(\frac{X_{i,t}}{V_{i,t}}\right) \quad \textit{for all } i. \tag{3.15}$$

Moreover, under the maintained non-degeneracy condition (i), conditions (ii) and (iii) are jointly necessary and sufficient *for the existence of a representation of the form* (3.15).

If, in addition, $g(\cdot)$ is continuous on $\mathcal{S}_t$ and the function $\varphi_t(\cdot)$ is linear in its argument, then there exists a scalar κ_t such that

$$\frac{g(Z_{i,t})}{V_{i,t}} = \kappa_t \frac{X_{i,t}}{V_{i,t}} \quad \textit{for all } i, \tag{3.16}$$

and $X_{i,t}/V_{i,t}$ is a unique sufficient statistic for expected excess returns in the sense of Definition 3.1. Conversely, if (3.16) *holds, then conditions (ii) and (iii) are satisfied. Equivalently,* (3.16) *holds if and only if*

$$g(Z_{i,t}) = \kappa_t X_{i,t} \quad \textit{for all } i,$$

so that linear intensity sufficiency in returns is equivalent to a level-proportional cash-flow exposure.

Proof is in Appendix A.1.

Proposition 3 characterizes the identifying restrictions that are necessary and sufficient for emission intensity to be uniquely sufficient in equilibrium. Homogeneity ensures that normalized exposure can depend on emission intensity and other firm characteristics, while scale invariance restricts this dependence to operate exclusively through intensity. Economically, homogeneity arises from additivity of cash-flow exposure under partitioning, as discussed below. When these conditions hold, two firms with the same ratio $X_{i,t}/V_{i,t}$ must have identical return-relevant exposure $g(Z_{i,t})/V_{i,t}$

regardless of their absolute scale. When these conditions fail, the exposure function $g(Z_{i,t})$ cannot be summarized by emission intensity alone, and additional characteristics beyond $X_{i,t}/V_{i,t}$ must enter priced exposure. Crucially, these restrictions translate into two sharp cross-sectional tests: a spanning restriction, under which intensity should absorb the return-relevant information in emission levels, and an invariance restriction, under which the intensity premium should not vary systematically with firm scale. Both can be implemented directly in the data.

The uniqueness result additionally requires that the function $\varphi_t(\cdot)$ be linear. This restriction is not implied by homogeneity or scale invariance alone. Under homogeneity and scale invariance, normalized exposure depends on firm characteristics only through emission intensity, so that $g(Z_{i,t})/V_{i,t} = \varphi_t(X_{i,t}/V_{i,t})$. If φ_t is further linear, exposure is proportional to intensity, $g(Z_{i,t})/V_{i,t} = \kappa_t X_{i,t}/V_{i,t}$, which is equivalent to the level-proportional cash-flow exposure $g(Z_{i,t}) = \kappa_t X_{i,t}$. In this case, emission intensity is a unique sufficient statistic in the sense of Definition 3.1. When linearity fails but homogeneity and scale invariance still hold, emission intensity remains sufficient in the sense that exposure depends only on $X_{i,t}/V_{i,t}$, but the representation is no longer unique: nonlinear transformations of intensity generate equivalent pricing relations, and linear regressions on $X_{i,t}/V_{i,t}$ would be misspecified. The empirical tests in Section 6.2 reject this weaker condition as well: emission levels retain explanatory power conditional on intensity, which is inconsistent with $g(Z_{i,t})/V_{i,t}$ being any function of $X_{i,t}/V_{i,t}$ alone.

The proposition further clarifies what does and does not hinge on market completeness. Completeness affects the equilibrium *price* of risk λ_t in (3.14), but it does not impose the scale and homogeneity restrictions on $g(\cdot)$ needed for intensity sufficiency. These implications motivate the empirical sufficiency tests in Section 6.2, which evaluate the scale-invariance and spanning restrictions associated with Proposition 3.

Finally, my analysis focuses on environments in which climate-related shocks enter firm cash flows in dollar terms. If transition risk instead operates through survival or default-type thresholds that generate losses proportional to firm value, then the homogeneity and scale-invariance conditions required for intensity sufficiency may hold. In that case, emission intensity would naturally emerge as the relevant priced characteristic. This regime is accommodated as a boundary case in my framework; the empirical results indicate that it has not dominated pricing over the sample period studied.

Economic foundations of the sufficiency conditions Proposition 3 shows that emission intensity is a sufficient statistic for priced exposure only when the exposure function $g(\cdot)$ is homogeneous of degree one in $(X_{i,t}, V_{i,t})$ and satisfies scale invariance. These are mathematical restrictions on the exposure function. This subsection provides their economic interpretation and explains why they are unlikely to hold in general.

The key economic property underlying homogeneity is *additivity* of exposure. Consider a firm with emissions $X_{i,t}$ and value $V_{i,t}$ that is split into n identical entities, each with emissions $X_{i,t}/n$ and value $V_{i,t}/n$. Additivity requires that the total dollar cash-flow exposure to the climate shock be preserved:

$$n \cdot g\left(\frac{X_{i,t}}{n}, \frac{V_{i,t}}{n}, Z_{-X,V}\right) = g(X_{i,t}, V_{i,t}, Z_{-X,V}), \tag{3.17}$$

where $Z_{-X,V}$ collects all other firm characteristics held fixed in the partition experiment.

To see why additivity implies homogeneity of degree one, consider two operations. First, splitting a firm into n identical entities implies

$$g\left(\frac{X_{i,t}}{n}, \frac{V_{i,t}}{n}\right) = \frac{g(X_{i,t}, V_{i,t})}{n}.$$

Second, merging n identical firms with emissions X and value V into a single firm with emissions nX and value nV implies

$$g(nX, nV) = ng(X, V).$$

Together, these operations yield $g(\lambda X, \lambda V) = \lambda g(X, V)$ for all positive rational λ. Under continuity, the relation extends to all $\lambda > 0$, establishing homogeneity of degree one in $(X_{i,t}, V_{i,t})$. Consequently, normalized exposure $g(Z_{i,t})/V_{i,t}$ depends on firm characteristics only through emission intensity $X_{i,t}/V_{i,t}$.

Additivity holds when climate shocks translate into cash-flow impacts that are mechanically proportional to emissions and independent of firm scale. A uniform carbon tax is the canonical example: a change $\Delta\tau$ generates a compliance cost $\Delta\tau\, X_{i,t}$, and splitting the firm preserves total exposure. This structure is also implicit in sustainable-discount-rate representations, in which expected returns take the form $r^x_{i,t} \approx r_t + (S_t - \tau_t)X_{i,t}/V_{i,t}$. In this case, the underlying exposure is $g(Z_{i,t}) = \eta X_{i,t}$, which satisfies additivity by construction.

This property can fail through two economically distinct channels. First, many real-world climate policies do not operate on a per-ton basis. Regulatory thresholds, disclosure rules, and

benchmark-based allocations often create fixed or level-based compliance costs that depend on absolute scale rather than proportional emissions. For example, firms above regulatory cutoffs must purchase allowances or install monitoring infrastructure, while smaller firms may avoid these costs even when aggregate emissions are unchanged. Similarly, net-zero benchmarks are typically allocated using sectoral budgets, installed capacity, or historical baselines rather than firm value, so two firms with identical emission intensity but different scale can face very different compliance gaps. In such settings, corporate structure affects total exposure and additivity fails at the level of the policy instrument.

Second, additivity can fail even under a perfectly uniform carbon tax. While the policy shock is linear per ton, its cash-flow impact depends on firms' pricing power, input substitution possibilities, contractual structure, and adjustment costs. Firms differ in their ability to pass carbon costs through to output prices, implying that identical emissions can translate into very different cash-flow effects. Because these pass-through rates vary systematically with market structure and technology, the mapping from carbon price shocks to cash flows is generically firm-specific and scale-dependent. Achieving additivity therefore requires not only a per-unit instrument but also homogeneous economic incidence across firms.

Together, these arguments imply that the homogeneity and scale-invariance restrictions require a joint set of conditions: climate exposure must operate through per-unit instruments, the incidence of those instruments must be homogeneous across firms, and regulation must avoid thresholds, benchmarks, or trajectory-based constraints. These conditions are restrictive. Relaxing any one of them causes exposure to depend on firm characteristics beyond emission intensity.

More broadly, the exposure function $g(Z_{i,t})$ is generically multidimensional. Emissions determine the direct incidence of per-unit policy shocks, but the cash-flow impact of a given shock is mediated by firm-specific characteristics—pricing power, demand elasticity, contractual rigidity, input substitution, and adjustment costs—that are not summarized by any single emission metric. The corresponding single-statistic benchmark, in which $g(Z_{i,t}) = \eta X_{i,t}$, should therefore be viewed as a tractable specification that isolates the direct policy channel by shutting down mediating heterogeneity, rather than as a description of the typical firm environment.

This logic applies symmetrically to any candidate emission characteristic, not only to intensity. Firms with identical emissions, identical intensity, and identical policy exposure can nevertheless have different cash-flow sensitivities to climate shocks when incidence is heterogeneous. Identifying the full exposure function $g(Z_{i,t})$—incorporating incidence heterogeneity, pass-through, and

technology—remains an open structural question to which this paper's selection framework provides the analytical foundation.

4 Climate Risk, Emission Intensity, and Priced Characteristics

In this section I embed emission-based climate metrics in a general linear-factor pricing framework. The key distinction is between (i) cash-flow exposures $g_k(Z_{i,t})$, which describe how climate-related shocks enter dividends, and (ii) return exposures, which are value-normalized and therefore load on $g_k(Z_{i,t})/V_{i,t}$. The central question is thus when the normalized exposure can be summarized by emission intensity $X_{i,t}/V_{i,t}$ alone. Proposition 3 characterizes the stringent conditions for a one-factor environment. The remainder of this section extends the representation to multiple factors and characterizes the implications for emission-based measures of climate risk.[7]

4.1 Climate-related risk as stochastic exposure to policy or technology shocks

Consider climate-related risk, with transition risk as a leading example, defined as stochastic shocks to climate policy, carbon prices, or technology that affect firms' cash flows. To fix ideas, let τ_{t+1} denote a random carbon price at date $t+1$ and suppose that, holding quantities fixed at date t, firm i's date-$t+1$ dividend contains a term

$$D_{i,t+1} = \tilde{D}_{i,t+1} - \tau_{t+1} X_{i,t}, \tag{4.1}$$

where $\tilde{D}_{i,t+1}$ collects all other components orthogonal (conditionally) to $\hat{\tau}_{t+1}$. Let $\hat{\tau}_{t+1} \equiv \tau_{t+1} - \mathbb{E}_t[\tau_{t+1}]$ denote the innovation in the carbon price and write

$$D_{i,t+1} = \mu_{i,t+1} + g_{i,t}^{(\tau)} \hat{\tau}_{t+1}, \qquad g_{i,t}^{(\tau)} = -X_{i,t},$$

with $\mu_{i,t+1}$ conditionally mean-independent of $\hat{\tau}_{t+1}$. In any equilibrium, expected excess returns satisfy

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = -\frac{\mathrm{Cov}_t(m_{t+1}, D_{i,t+1})}{\mathbb{E}_t[m_{t+1}]\, V_{i,t}},$$

[7] I now adopt time subscripts. All static results in Sections 2-3 carry over to the dynamic setting under standard CARA-normal assumptions (see Internet Appendix D).

where $V_{i,t}$ denotes the date-t market value (price) of the claim to $D_{i,t+1}$. Substituting the dividend representation into the covariance pricing identity yields

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \lambda_t^{(\tau)} \frac{g_{i,t}^{(\tau)}}{V_{i,t}} = -\lambda_t^{(\tau)} \frac{X_{i,t}}{V_{i,t}}, \tag{4.2}$$

where the price of transition risk associated with carbon-price shocks is defined by

$$\lambda_t^{(\tau)} = -\frac{\mathrm{Cov}_t(m_{t+1}, \hat{\tau}_{t+1})}{\mathbb{E}_t[m_{t+1}]}.$$

In the special case of a carbon-price shock with linear incidence, exposure is proportional to emissions, but in general the exposure functions $g_k(Z_{i,t})$ need not collapse to any single emission-based measure.

More generally, allowing for multiple climate-related channels, dividends take the form

$$D_{i,t+1} = \tilde{D}_{i,t+1} + \sum_{k=1}^{K} g_k(Z_{i,t}) f_{k,t+1}, \tag{4.3}$$

where $f_{k,t+1}$ are mean-zero climate-related factors (e.g., policy, technology, or demand shocks) and $g_k(Z_{i,t})$ are dollar exposure coefficients that depend on characteristics $Z_{i,t}$, which may include $X_{i,t}$, $V_{i,t}$, $X_{i,t}/V_{i,t}$, changes in emissions, efficiency measures, and so on.

4.2 Market completeness, prices of risk, and exposure functions

Building on this distinction, this subsection examines how market completeness affects equilibrium pricing. Complete versus incomplete markets affect the prices of climate-related risk but not, for a given technology and shock structure, the mechanism-implied dollar exposure functions $g_k(Z_{i,t})$. To make this precise, consider the following general linear-factor environment.

Let $\{f_{k,t+1}\}_{k=1}^{K}$ be conditionally mean-zero factors, and suppose that dividends satisfy (4.3):

$$D_{i,t+1} = \mu_{i,t+1} + \sum_{k=1}^{K} g_k(Z_{i,t}) f_{k,t+1},$$

with $\mu_{i,t+1}$ conditionally independent of $\{f_{k,t+1}\}_k$ and $Z_{i,t}$ $\mathcal{F}_t$–measurable. Assume that the stochas-

tic discount factor admits the linear projection representation

$$m_{t+1} = a_t + \sum_{k=1}^{K} b_{k,t} f_{k,t+1} + \varepsilon_{t+1}, \tag{4.4}$$

where $a_t > 0$, $\{b_{k,t}\}_k$ are $\mathcal{F}_t$–measurable, and ε_{t+1} is conditionally mean-zero and orthogonal to all factors and traded payoffs. Under complete markets, the SDF is unique and coincides with the representative agent's or planner's marginal rate of substitution. Under incomplete markets, the SDF is not unique, but admits a representation of the form (4.4), and in our CARA-normal environment equilibrium pins down a particular affine specification.

Using the Euler equation $\mathbb{E}_t[m_{t+1}R_{i,t+1}] = 1$ together with (4.3) and (4.4) yields the following characterization.

Proposition 4. *Suppose that dividends are given by (4.3) and that the SDF admits the representation (4.4), with $\{f_{k,t+1}\}_k$ mutually uncorrelated conditional on $\mathcal{F}_t$. Then expected excess returns satisfy*

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \sum_{k=1}^{K} \lambda_{k,t} \frac{g_k(Z_{i,t})}{V_{i,t}}, \tag{4.5}$$

where

$$\lambda_{k,t} = -\frac{b_{k,t} \operatorname{Var}_t(f_{k,t+1})}{\mathbb{E}_t[m_{t+1}]}, \qquad k = 1, \dots, K, \tag{4.6}$$

are the factor prices of risk. In particular, for each factor k, the priced return exposure is the value-normalized object $g_k(Z_{i,t})/V_{i,t}$: two assets with identical normalized exposures for all k have identical expected excess returns, regardless of how $g_k(\cdot)$ depends on primitives such as $(X_{i,t}, V_{i,t}, \dots)$.

Proof. Substituting (4.3) and (4.4) into the covariance pricing identity

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = -\frac{\operatorname{Cov}_t(m_{t+1}, D_{i,t+1})}{\mathbb{E}_t[m_{t+1}]\, V_{i,t}}$$

and using orthogonality of ε_{t+1} and conditional mutual uncorrelatedness of the factors yields

$$\operatorname{Cov}_t(m_{t+1}, D_{i,t+1}) = \sum_{k=1}^{K} b_{k,t}\, g_k(Z_{i,t}) \operatorname{Var}_t(f_{k,t+1}),$$

which implies (4.5) and (4.6). □

Proposition 4 holds irrespective of whether markets are complete or incomplete. Market completeness affects the coefficients $b_{k,t}k$ and thus prices of risk $\lambda k, tk$, but not the exposure functions

$g_k(Z_{i,t})$.

The separation between market structure and the exposure function $g_k(\cdot)$ is exact when firm characteristics are taken as predetermined at date t. In richer general-equilibrium environments, firms may adjust emissions, technology, or production in response to expected policy and risk premia. In such settings, the equilibrium mapping $g_k(Z_{i,t})$ can itself depend on the price of risk through firms' optimal decisions. Proposition 4 should therefore be interpreted as characterizing the pricing of exposures conditional on the characteristics realized at date t. The dynamic production model in Appendix D provides an example in which exposures are endogenous, yet the normalized representation (4.5) continues to describe the cross section of expected returns.

In the transition-tax example (4.1) with a single factor $f_{1,t+1} = \hat{\tau}_{t+1}$, Proposition 4 identifies $g_1(Z_{i,t})$ with $-X_{i,t}$ and thus the priced return exposure with $-X_{i,t}/V_{i,t}$. More generally, when $g_k(Z_{i,t})$ depends on emission levels, changes, technology, or policy allocations, the resulting priced characteristics are the corresponding normalized exposure functions $g_k(Z_{i,t})/V_{i,t}$. The next section illustrates how alternative specifications of $g_k(\cdot)$ map into distinct empirical predictors.

5 Illustrative Exposure Specifications

Section 3.4 showed that emission intensity is a sufficient statistic only under economically demanding homogeneity and scale-invariance restrictions. This section complements that result by deriving exposure functions from economically motivated mechanisms. The incomplete-markets equilibrium in Section 3 relied on a benchmark in which dividends load linearly on emission levels. Realistic climate exposure, however, is richer. Policy shocks need not scale only with current emissions; they can depend on recent changes, distances from regulatory benchmarks, and technology choices that shape how cash flows respond to climate risk. The three specifications that follow–transition losses proportional to levels (the Section 3 benchmark), exposures depending jointly on levels and changes (motivated by trajectory-based regulation), and path-based exposures under net-zero commitments–illustrate a general lesson: once exposure functions are disciplined by realistic policy design and technology, priced characteristics need not collapse to any single emission-based statistic.

5.1 Exposure proportional to emission levels

I start with the simplest case: cash-flow exposure proportional to emission levels. For example, if a per-ton carbon price increases by \$10 and a firm emits 100 kilotons, its compliance cost rises by \$1 million. A reduced-form representation of this channel is $\theta_{i,t} = \eta X_{i,t}$, linear in emissions.[8] This specification is natural when policy operates through per-unit instruments–carbon taxes, permit obligations, or output restrictions that scale directly with total emissions.

Under $\theta_{i,t} = \eta X_{i,t}$, the priced exposure is therefore $\eta X_{i,t}/V_{i,t}$, so emission intensity emerges mechanically as a sufficient statistic for pricing. This sufficiency, however, is not generic: it holds only because the primitive cash-flow channel is linear in levels and does not depend on other scale objects. Once exposure depends on additional characteristics such as changes in emissions, compliance constraints, or distance from a regulatory path, two firms with the same $X_{i,t}/V_{i,t}$ can have different $g(Z_{i,t})/V_{i,t}$, and intensity no longer spans priced exposure.

Single-factor setup. To formalize this, consider a single climate factor ζ_{t+1} with $\mathbb{E}_t[\zeta_{t+1}] = 0$ and write firm i's dividend as

$$D_{i,t+1} = \tilde{D}_{i,t+1} + \theta_{i,t}\,\zeta_{t+1}, \tag{5.1}$$

where $\tilde{D}_{i,t+1}$ is orthogonal to ζ_{t+1}. The cash-flow exposure satisfies

$$\theta_{i,t} = g(Z_{i,t}), \tag{5.2}$$

where $g(\cdot)$ maps firm characteristics into dollar sensitivities. In the Section 3 benchmark, $\theta_{i,t} = \eta X_{i,t}$.

Expected excess returns then satisfy

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \lambda_t \frac{g(Z_{i,t})}{V_{i,t}}. \tag{5.3}$$

From mechanisms to empirical predictors. The functional form of $g(\cdot)$ is disciplined by economic structure: it reflects how climate shocks enter cash flows. If $g(Z_{i,t}) = \eta X_{i,t}$, emission intensity emerges as the return exposure. But if policy operates through adjustment costs, compliance constraints, or net-zero benchmarks, $g(Z_{i,t})$ naturally depends on additional characteristics, such as emission changes $\Delta X_{i,t}$ or distances from regulatory paths, and the resulting exposure need

[8] To streamline notation, I denote the dollar exposure to the climate factor by $\theta_{i,t}$.

not collapse to intensity.

The theory thus disciplines measurement: an emission-based predictor is admissible only if it can be derived from a coherent cash-flow mechanism, and it is empirically sufficient if it spans the return variation associated with $g(Z_{i,t})/V_{i,t}$. Intensity is sufficient only when dollar exposures satisfy the homogeneity and scale-invariance conditions in Proposition 3–conditions that fail once $g(\cdot)$ depends on characteristics beyond $(X_{i,t}, V_{i,t})$.

Multi-factor extension. With multiple climate shocks $\boldsymbol{\zeta}_{t+1} = (\zeta_{1,t+1}, \ldots, \zeta_{K,t+1})^\top$, dividends become

$$D_{i,t+1} = \tilde{D}_{i,t+1} + \sum_{k=1}^{K} \theta_{i,t}^{(k)} \zeta_{k,t+1}, \qquad \theta_{i,t}^{(k)} = g_k(Z_{i,t}),$$

and expected returns inherit the structure

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \sum_{k=1}^{K} \lambda_t^{(k)} \frac{g_k(Z_{i,t})}{V_{i,t}}. \tag{5.4}$$

Each factor $\zeta_{k,t+1}$ is priced through its exposure function $g_k(Z_{i,t})/V_{i,t}$. Physical damage shocks, carbon-price innovations, and technology shocks can each generate distinct exposure functions, and there is no reason to expect a single characteristic such as intensity to span all of them.

Invariance restrictions on exposure functions. Not every function $g(\cdot)$ is economically coherent. Dollar exposures must satisfy basic invariance requirements. Splitting a firm into two identical copies with half the emissions and half the value should preserve total exposure (additivity). Holding technology fixed, higher emissions should not reduce exposure to a per-unit tax or damage shock (monotonicity). These restrictions ensure that exposure functions correspond to actual economic mechanisms rather than arbitrary functional forms.

Under these disciplines, emission levels, intensities, changes, and SCC-weighted measures are all admissible–each can be derived from a specific policy or technology channel. But they are empirically sufficient only if they span the priced exposure $g_k(Z_{i,t})/V_{i,t}$. When cash flows depend on characteristics beyond those captured by a given metric, that metric fails the spanning tests in Section 6.2.

5.2 Transition losses depending on levels and changes

Policy tightening does not just penalize current emissions–it targets firms moving in the wrong direction. A coal plant ramping up capacity faces different regulatory risk than one phasing down, even if both have similar emission levels today. If enforcement responds to trajectories rather than stocks, cash-flow exposure depends on both $X_{i,t}$ and $\Delta X_{i,t}$, and intensity alone cannot span the priced risk.

To see this, consider a policy shock that arrives when regulators perceive insufficient progress toward transition goals. Firms increasing emissions signal resistance to decarbonization and face steeper penalties when enforcement tightens. Formally, let transition losses depend on both the current emission stock and recent changes:

$$\Gamma_{i,t} = \phi_1 X_{i,t} + \phi_2 \Delta X_{i,t}, \tag{5.5}$$

where $\Delta X_{i,t} = X_{i,t} - X_{i,t-1}$ and $\phi_1, \phi_2 > 0$. The level $X_{i,t}$ captures baseline exposure; the change $\Delta X_{i,t}$ captures trajectory risk.

Policy-tightening shocks and transition losses. Let $J_{t+1} \in \{0, 1\}$ denote a binary tightening event with $\mathbb{E}_t[J_{t+1}] = p_t$, and define the innovation

$$\zeta_{t+1} = J_{t+1} - p_t.$$

The shock J_{t+1} captures unanticipated regulatory acceleration–stricter enforcement, faster phase-outs, or abrupt carbon-price increases. Firm i's dividend becomes

$$D_{i,t+1} = \bar{D}_{i,t+1} - \Gamma_{i,t} J_{t+1}, \tag{5.6}$$

where $\bar{D}_{i,t+1}$ is the dividend absent tightening and $\Gamma_{i,t}$ is the dollar loss when tightening occurs. This loss reflects stranded assets, accelerated abatement costs, or compliance penalties that hit when policy shifts earlier than expected.

Substituting (5.5) into (5.6) and using $\zeta_{t+1} = J_{t+1} - p_t$ yields

$$D_{i,t+1} = \bar{D}_{i,t+1} - (\phi_1 X_{i,t} + \phi_2 \Delta X_{i,t}) p_t - (\phi_1 X_{i,t} + \phi_2 \Delta X_{i,t}) \zeta_{t+1}.$$

The cash-flow exposure to the transition factor is

$$g_\zeta(Z_{i,t}) = -(\phi_1 X_{i,t} + \phi_2 \Delta X_{i,t}),$$

where $Z_{i,t}$ includes both levels and changes.

Equilibrium pricing with trajectory dependence. Under the log-linear approximation in Proposition 4, expected excess returns satisfy

$$\mathbb{E}_t[R_{i,t+1} - R_{f,t}] = \lambda_{\zeta,t} \frac{\phi_1 X_{i,t} + \phi_2 \Delta X_{i,t}}{V_{i,t}} + \text{other factor terms.} \tag{5.7}$$

Both levels and changes enter through the normalized exposure $(\phi_1 X_{i,t} + \phi_2 \Delta X_{i,t})/V_{i,t}$. When $\phi_2 = 0$, the model collapses to the Section 5.1 baseline where exposure scales with levels alone. When $\phi_2 > 0$, firms ramping up emissions bear additional transition risk that cannot be captured by $X_{i,t}/V_{i,t}$ alone.

This has direct implications for intensity-based screens. Two firms with identical emission intensity can have different priced exposure if one is increasing emissions ($\Delta X_{i,t} > 0$) while the other is cutting ($\Delta X_{i,t} < 0$). The sufficiency conditions in Proposition 3 fail because the exposure function depends on an additional characteristic–recent emission changes–that is not summarized by the intensity ratio $X_{i,t}/V_{i,t}$.

Empirical implications. Equation (5.7) suggests that the appropriate empirical specification includes both $X_{i,t}/V_{i,t}$ and $\Delta X_{i,t}/V_{i,t}$ as predictors. A regression of returns on these characteristics is a reduced-form test of whether trajectory-based exposure is priced. The coefficients need not recover (ϕ_1, ϕ_2) separately, as they combine technological sensitivities with the equilibrium price of risk $\lambda_{\zeta,t}$, but the joint specification is the natural empirical analogue of the mechanism in (5.5).

The incomplete-markets structure, heterogeneous climate concern, and SDF decomposition from Section 3 remain unchanged. What changes is the exposure function $g_\zeta(Z_{i,t})$, which now depends on both stocks and flows. This illustrates the general point: different policy designs generate different priced characteristics, and intensity suffices only under the knife-edge restriction that exposure depends solely on $(X_{i,t}, V_{i,t})$.

5.3 Net-zero objectives and path-based exposure

Net-zero commitments change the regulatory game. A firm emitting 150 kilotons with a 100-kiloton benchmark faces different transition risk than a firm at 75 kilotons with a 50-kiloton benchmark–even though both are 50% above their targets. What matters is not just the intensity ratio but the absolute distance from a mandated path, and that distance varies with scale even when $X_{i,t}/V_{i,t}$ is fixed.

This matters because net-zero benchmarks are typically allocated using level-based rules–sectoral budgets, capacity shares, historical baselines–not intensity targets. Two firms with identical emission intensity can have vastly different compliance trajectories if one is large and distant from its benchmark while the other is small and close. If enforcement tightens or phase-out schedules accelerate, exposure scales with the gap $(X_{i,t} - X_{i,t}^{NZ})$, not just with current emissions or their ratio to firm value. Because $X_{i,t}^{NZ}$ is not proportional to $V_{i,t}$ under level-based allocations, the normalized gap $(X_{i,t} - X_{i,t}^{NZ})/V_{i,t}$ varies across firms even when $X_{i,t}/V_{i,t}$ is held fixed. Intensity alone therefore cannot span priced exposure in (5.9).

Path-based exposure under net-zero regulation. Let $X_{i,t}^{NZ}$ denote firm i's net-zero-consistent benchmark at date t, constructed from sectoral pathways, regulatory allocations, or firm commitments. This benchmark is predetermined at t and reflects policy design or historical negotiations rather than equilibrium outcomes; it may be measured with error in practice. Let ζ_{t+1} denote net-zero transition news: an acceleration of phase-out schedules, stricter enforcement, or increases in marginal abatement costs along the transition path.

Cash-flow exposure to ζ_{t+1} reflects two channels. First, per-unit losses that scale with current emissions $X_{i,t}$ (permit obligations, tax exposure). Second, repricing and compliance risk tied to the distance from the benchmark, $(X_{i,t} - X_{i,t}^{NZ})$. Formally, let

$$g_\zeta(Z_{i,t}) = \vartheta_1 X_{i,t} + \vartheta_2 (X_{i,t} - X_{i,t}^{NZ}), \tag{5.8}$$

where $\vartheta_1, \vartheta_2 \geq 0$. The level $X_{i,t}$ captures baseline exposure, while the gap $(X_{i,t} - X_{i,t}^{NZ})$ captures path-deviation risk under net-zero regulation.[9]

[9] One can additionally allow adjustment-speed penalties proportional to $\Delta X_{i,t}$ as in Section 5.2.

Equilibrium pricing with path dependence. Under the log-linear approximation in Proposition 4, expected excess returns satisfy

$$\mathbb{E}_t[R_{i,t+1} - R_{f,t}] = \lambda_{\zeta,t} \frac{g_\zeta(Z_{i,t})}{V_{i,t}} + \text{other factor terms.} \tag{5.9}$$

The priced exposure is the normalized function

$$\frac{g_\zeta(Z_{i,t})}{V_{i,t}} = \vartheta_1 \frac{X_{i,t}}{V_{i,t}} + \vartheta_2 \frac{X_{i,t} - X_{i,t}^{NZ}}{V_{i,t}}.$$

Intensity $X_{i,t}/V_{i,t}$ alone cannot span this exposure. Even firms with identical intensity can have sharply different path deviations $(X_{i,t} - X_{i,t}^{NZ})/V_{i,t}$ because benchmarks are allocated by levels or capacity shares, not intensity ratios. A large firm at 100% above its benchmark carries fundamentally different compliance risk than a small firm at 100% above its benchmark—the absolute emission gap that must be closed differs by scale.

The sufficiency conditions in Proposition 3 require that priced exposure depend only on $(X_{i,t}, V_{i,t})$ and satisfy homogeneity and scale invariance. Net-zero exposures violate these conditions in two ways. First, they depend on an additional predetermined characteristic, $X_{i,t}^{NZ}$, that reflects policy design rather than current firm decisions. Second, because benchmarks are level-based, the gap $(X_{i,t} - X_{i,t}^{NZ})/V_{i,t}$ varies with scale even when intensity is fixed. Consequently, intensity-based screens systematically misclassify firms when compliance risk is path-dependent.

Equation (5.9) suggests that empirical specifications motivated by path-based regulation include not only $X_{i,t}/V_{i,t}$ but also the normalized deviation $(X_{i,t} - X_{i,t}^{NZ})/V_{i,t}$. A spanning test that includes only intensity implicitly imposes $\vartheta_2 = 0$, that is, that transition risk is insensitive to distance from the benchmark. This restriction is economically strong in settings where compliance is assessed relative to net-zero pathways rather than current flow emissions alone.

Net-zero objectives thus provide a concrete illustration of the paper's central message: the appropriate emission-based predictor is determined by policy design, not by intensity a priori. When regulation operates through path constraints, the priced characteristic is the exposure function $g_\zeta(Z_{i,t})/V_{i,t}$, which depends on levels, gaps, and changes jointly. Intensity sufficiency is a restrictive special case that holds only when transition risk ignores path deviations and adjustment dynamics–conditions unlikely to hold in environments where cumulative emissions or trajectory compliance drives enforcement.

To conclude this section, Table 1 summarizes the model's implications for empirical implemen-

tation. For each economic mechanism, formalized through an exposure function $g(Z_{i,t})$, the table reports the object that prices equilibrium returns and the corresponding variable implied for return regressions. The table serves as a compact reference, abstracting from regression controls and normalization that are common across specifications.

6 An Application to Emission Intensity

This section applies the sufficiency framework developed in Sections 3–5 to emission intensity, the leading policy-relevant candidate characteristic for priced climate exposure. Section 6.1 draws out the implications of the framework for interpreting emission-based measures of climate risk and states the two cross-sectional restrictions–spanning and scale invariance–that must hold if intensity is a sufficient statistic under Proposition 3. Section 6.2 tests these restrictions directly using an international panel of firms and reports that both are rejected, with the rejection strongest among smaller firms. The analysis does not adjudicate between alternative emission-based proxies or attempt to estimate a structural exposure function; these are questions for future research.

6.1 Interpreting intensity-based measures

This subsection draws out the empirical implications of the framework for interpreting emission-based measures of climate risk. The central distinction is between structural exposure functions $g(Z_{i,t})$, which govern cash-flow sensitivities, and emission-based characteristics such as $X_{i,t}/V_{i,t}$ that are used to approximate them. In equilibrium, expected excess returns load on the exposure $g(Z_{i,t})/V_{i,t}$. Any scalar characteristic is therefore informative for pricing only to the extent that it recovers this underlying object, and commonly used measures should be interpreted as projections of $g(Z_{i,t})/V_{i,t}$ onto selected firm attributes rather than as primitives of the pricing relation.

Intensity-based approaches implicitly impose the homogeneity and scale-invariance conditions required for intensity sufficiency, characterized in Proposition 3. When these conditions hold, priced exposure depends on firm characteristics only through $X_{i,t}/V_{i,t}$, and intensity is a valid scalar summary. When they fail, priced characteristics reflect the underlying exposure functions determined by technology, regulation, and firm behavior, and intensity-based measurement misses the component of exposure that is not summarized by the ratio. Section 5 shows that these functions may depend on emission levels (when per-unit policies scale losses with $X_{i,t}$), emission changes (when regulation responds to trajectories), net-zero gaps (when compliance is path-dependent), or

efficiency choices (when technology affects shock sensitivity). The appropriate predictor depends on which economic mechanism governs the mapping from firm characteristics into priced cash-flow exposure.

Viewed through this lens, the conflicting findings in the empirical literature—with some studies documenting premia on levels, others on intensity, and still others on changes or composite scores—are less puzzling than they may appear. Different empirical measures correspond to projections of $g(Z_{i,t})/V_{i,t}$ onto different firm attributes, and the relative informativeness of each projection depends on which economic channels are active in the sample being studied. Differences across studies therefore reflect differences in the underlying exposure channels being captured, not competing proxies for a common carbon premium.

Proposition 3 yields two cross-sectional restrictions that must hold if emission intensity is a sufficient statistic for priced climate exposure. First, *scale invariance*: the premium per unit of $X_{i,t}/V_{i,t}$ must not vary systematically with firm size. This restriction is tested in Table 3 by allowing the intensity premium to differ across size bins and testing the implied cross-bin equality. Second, *spanning*: once firm value is controlled for, emission levels should carry no incremental explanatory power for expected returns. Because $\log(X_{i,t}/V_{i,t}) = \log X_{i,t} - \log V_{i,t}$, this restriction is equivalent to the linear coefficient restriction $b_{X,t} + b_{V,t} = 0$ in a regression with $\log X_{i,t}$ and $\log V_{i,t}$ entered separately, and it is tested in Tables 4 and 5. Because Proposition 3 establishes that conditions (ii) and (iii) are jointly necessary and sufficient for intensity sufficiency, both restrictions are equivalent implications of sufficiency, and rejection of either suffices to reject intensity sufficiency. The spanning test evaluates whether $g(Z_{i,t})/V_{i,t}$ depends on firm characteristics only through $X_{i,t}/V_{i,t}$, so its rejection rules out intensity sufficiency in both the linear form of equation (3.16) and the weaker nonlinear form in which $g(Z_{i,t})/V_{i,t} = \phi_t(X_{i,t}/V_{i,t})$ for some function ϕ_t. This is inconsistent with the empirical finding that emission levels retain explanatory power conditional on firm value.

These tests are not tests of market completeness per se. As shown in Section 5, market structure affects the equilibrium price of risk λ_t but does not impose the scale and homogeneity restrictions on $g(\cdot)$ that make intensity sufficient. The restrictions tested below are therefore properties of the exposure function itself, not of the pricing kernel. Their rejection clarifies that the debate over emission-based measures reflects distinct economic questions—policy implementation versus equilibrium risk measurement—rather than alternative representations of the same priced object.[10]

[10] A formal welfare analysis is provided in Appendix C, which characterizes the consumption-equivalent loss from incomplete markets and clarifies its distinction from equilibrium carbon premia.

6.2 Empirical evidence on sufficiency

I test the scale-invariance and spanning restrictions implied by Proposition 3 using an international panel of firms. These restrictions imply that, if emission intensity is sufficient, the pricing of intensity must be independent of firm scale and emission levels must carry no incremental explanatory power conditional on intensity. The objective is to evaluate whether emission intensity, the leading policy-relevant candidate, satisfies the cross-sectional restrictions required for sufficiency under Proposition 3. It is not to identify a structural exposure function, to rank alternative predictors, or to advocate a preferred proxy. The tests below are restriction tests: under the null of intensity sufficiency, the restrictions must hold in the data regardless of which underlying mechanism generates climate exposure.

Under the equilibrium pricing framework in Section 3, cross-sectional differences in expected returns reflect differences in normalized exposure functions $g(Z_{i,t})/V_{i,t}$. Because firm characteristics are predetermined with respect to subsequent shocks, predictable return differences associated with these characteristics provide reduced-form evidence on the underlying exposure function. These tests are reduced-form and do not separately identify structural exposure from endogenous firm responses: because emissions and firm value are equilibrium objects, cross-sectional variation in $X_{i,t}$ and $V_{i,t}$ may reflect both underlying exposure and firms' optimal adjustments to expected climate and policy risk. The regressions below therefore test whether emission intensity alone summarizes priced exposure or whether additional scale-dependent components are required.

I implement a direct test of scale invariance by allowing the intensity premium to vary across size bins and testing the implied cross-bin equality restriction. I then conduct spanning tests that evaluate whether emission levels carry independent explanatory power conditional on firm value. Finally, I report within-size-bin spanning tests that address the concern that pooled results reflect residual slope heterogeneity associated with firm scale.

The tests use an international panel of 13,373 firms representing 92 countries and 438 Trucost industries over 2005–2022. I collect panel data on stock returns $R_{i,t+1}$, the risk-free rate $R_{f,t}$, firm-level Scope 1 emissions $X_{i,t}$, firm value $V_{i,t}$, and standard firm characteristics $W_{i,t}$ (such as market capitalization, market-to-book ratio, book leverage, profitability, investment, momentum, age, PPE, and volatility). The emissions data come from S&P Trucost, while all other information is from Capital IQ. Following Pedersen (2025), I define emission intensity as $X_{i,t}/V_{i,t}$, where $V_{i,t}$ denotes the market value of equity, and I define the percentage change in emissions as $\Delta X_{i,t} =$

$(X_{i,t} - X_{i,t-1})/X_{i,t-1}$. To mitigate the influence of outliers, I winsorize intensity and emissions growth at the 2.5% level. Table 2 reports summary statistics for the key empirical variables.

A direct test of scale invariance. A central implication of intensity sufficiency is a scale-invariance restriction: firms with the same intensity but different scale must have the same priced transition exposure. If this restriction holds, the premium per unit of intensity should not vary systematically with firm scale. To test it, each period I sort firms by size (market value of equity) into five equal-count bins and estimate the pooled regression

$$R_{i,t+1} - R_{f,t} \;=\; \alpha_t \;+\; \sum_{b=1}^{B} b_b \, D_{b,i,t} \log\left(\tfrac{X_{i,t}}{V_{i,t}}\right) + \gamma_t^{\top} W_{i,t} \;+\; \varepsilon_{i,t+1}, \tag{6.1}$$

where $D_{b,i,t}$ indicates membership in size bin b. The specification includes $W_{i,t}$, industry $\times$ time fixed effects, and country fixed effects. Standard errors are double-clustered by industry and time. The null implied by Proposition 3 is $H_0 : b_1 = \cdots = b_B$, which I test using a Wald statistic based on the clustered covariance matrix.

Table 3 reports the results. Column (1) presents the baseline specification with size-bin interactions only; column (2) augments the specification with the full set of firm-level controls. In both columns, the estimated intensity premium varies materially across size bins: the coefficient on $\log(X/V)$ for the smallest firms is between 0.176 and 0.208, and interactions with larger bins enter monotonically negatively, with coefficients ranging from -0.036 to -0.064. The intensity premium is therefore largest among the smallest firms and declines monotonically with firm size. The Wald tests of the cross-bin equality restriction, reported at the foot of the table, reject the null at the 5% level in both specifications ($p = 0.049$ and $p = 0.023$, respectively). This rejection provides direct evidence against the hypothesis that emission intensity summarizes priced transition exposure in a scale-invariant manner.

Spanning tests: does intensity subsume levels? I complement the scale-invariance tests with spanning tests that evaluate whether emission levels carry independent explanatory power once firm value is controlled for. Because $\log(X_{i,t}/V_{i,t}) = \log X_{i,t} - \log V_{i,t}$, intensity sufficiency can be recast as a linear coefficient restriction on a regression that includes $\log X_{i,t}$ and $\log V_{i,t}$ separately: intensity enters returns as a sufficient statistic if and only if the coefficients on the two

components are equal and opposite. I therefore estimate the pooled regression

$$R_{i,t+1} - R_{f,t} \;=\; \alpha_t \;+\; b_{X,t} \log X_{i,t} \;+\; b_{V,t} \log V_{i,t} \;+\; \gamma_t^\top W_{i,t} \;+\; \varepsilon_{i,t+1}, \tag{6.2}$$

and test the restriction $H_0 : b_{X,t} + b_{V,t} = 0$ using a Wald statistic based on the clustered covariance matrix.

Table 4 reports the results. Emission levels enter with positive and statistically significant coefficients at both the one- and twelve-month horizons, while $\log V_{i,t}$ enters negatively and significantly. The Wald test of $b_{X,t} + b_{V,t} = 0$ is rejected at the 1% level in the one-month specification (column 1; $p = 0.008$) and at the 5% level at the twelve-month horizon (column 4; $p = 0.031$). These estimates indicate that emission levels carry independent cross-sectional information about expected returns beyond what is captured by the ratio $X_{i,t}/V_{i,t}$, which is inconsistent with intensity sufficiency. Columns (2)–(3) and (5)–(6) report analogous tests using sales as an alternative scaling variable and confirm that the spanning rejection is not specific to the market-value normalization. The rejection of the spanning restriction is not driven by small firms or extreme values of the intensity ratio. Re-estimating the pooled specification after excluding the lowest market-cap quintile yields similar results: the null hypothesis $b_{X,t} + b_{V,t} = 0$ continues to be rejected at conventional significance levels. Because small firms are most likely to exhibit measurement error in emissions and extreme values of $X_{i,t}/V_{i,t}$, this result indicates that the spanning rejection is not an artifact of denominator-driven noise or illiquidity concentrated in the lower tail of the size distribution, but reflects a broader feature of the cross section.

These findings have a direct implication for the underlying exposure function. From the equilibrium pricing representation in Section 3,

$$\mathbb{E}_t[R_{i,t+1} - R_{f,t}] \;\approx\; \lambda_t \, \frac{g(Z_{i,t})}{V_{i,t}},$$

cross-sectional variation in expected returns reflects differences in normalized cash-flow exposure. If expected returns vary with emission levels conditional on firm value, then firms with identical $X_{i,t}/V_{i,t}$ but different scale must have different values of $g(Z_{i,t})/V_{i,t}$. This implies that dollar exposure $g(Z_{i,t})$ cannot be proportional to emissions alone and instead must scale non-proportionally with firm value, so that normalized exposure depends on firm scale. The data are therefore consistent with mechanisms such as fixed compliance costs, scale-dependent pass-through, regulatory thresholds, or adjustment frictions, which generate exposure functions that depend separately on

emission levels and firm scale.

The restrictions being tested are economically meaningful. In column (1) of Table 4, a one-standard-deviation increase in log(Scope 1) (approximately 2.85; see Table 2) predicts a 0.36% increase in next-month expected returns, or roughly 4.3% annualized. At the twelve-month horizon (column 4), the corresponding effect is 0.15% per month, or about 1.8% annualized. Thus, firms with similar $X_{i,t}/V_{i,t}$ but different emission scales earn systematically different expected returns, directly contradicting the scale-invariance and spanning restrictions implied by Proposition 3.

Within-size-bin spanning. A remaining concern is that the spanning asymmetry in Table 4 reflects slope heterogeneity associated with firm scale. If the price of transition exposure varies with size, then emission levels may appear to add explanatory power in pooled regressions even when intensity is the relevant exposure measure within scale classes. To address this concern directly, I re-estimate the joint levels–value specification (6.2) separately within size quintiles formed using market value of equity, using the same controls and fixed effects. By construction, firm scale is approximately fixed within each regression, so this design isolates whether emission levels add return-relevant information beyond firm value even among firms of comparable size.

Table 5 reports the results. Emission levels enter with positive and statistically significant coefficients across all five size quintiles, with estimates ranging from 0.100 to 0.164 and significance at the 5% level or better. The Wald test of the homogeneity restriction $b_{X,t} + b_{V,t} = 0$ is rejected at the 1% level in the smallest quintile ($p = 0.004$) and marginally rejected in the next two ($p = 0.075$ and $p = 0.089$). In the two largest quintiles, by contrast, the restriction cannot be rejected at conventional levels ($p = 0.238$ and $p = 0.326$). The monotone pattern of coefficients and p-values across bins—largest intensity premium and sharpest rejection among the smallest firms, progressive weakening with scale—is informative about where intensity-based measurement breaks down most clearly, without requiring commitment to any specific mechanism.[11] What matters for the tests reported here is that the pooled rejection of intensity sufficiency in Tables 3 and 4 is not an artifact of residual size heterogeneity, since levels retain independent explanatory power throughout the cross-section and significantly so across three of the five size quintiles.

Concerns that the joint specification merely reflects mechanical ratio relationships are mitigated

[11] The monotone pattern across size bins is also consistent with microstructure effects concentrated in smaller firms, such as illiquidity or bid-ask bounce. Two features of the evidence mitigate this concern. First, emission levels retain statistically significant explanatory power within each size quintile, including the largest firms where microstructure concerns are minimal. Second, the specifications include controls for volatility and momentum that absorb standard microstructure proxies.

by both the correlation structure of the data and the within-size-bin evidence. Emission intensity and log emissions are only moderately correlated, and intensity is close to orthogonal to firm size. Moreover, emission levels retain independent explanatory power even within narrow size bins, where firm scale is approximately fixed. These results indicate that the spanning rejection reflects genuine differences in priced exposure rather than mechanical denominator effects arising from the definition of intensity.

The rejection does not rely on long-horizon predictability or on the magnitude of estimated premia. The objects being tested are cross-sectional restrictions implied by Proposition 3—scale invariance and spanning–which must hold at any horizon if emission intensity is a uniquely sufficient statistic for priced exposure.

Changes in emissions. To connect more directly to the transition-risk channel in Section 5, I also consider a specification that allows both the level and the recent change in emissions to enter the exposure function. Specifically, I estimate

$$R_{i,t+1} - R_{f,t} \;=\; \alpha_t \;+\; b_{X,t} \log X_{i,t} \;+\; b_{\Delta X,t}\, \Delta X_{i,t} \;+\; \gamma_t^{\top} W_{i,t} \;+\; \varepsilon_{i,t+1}, \tag{6.3}$$

and report the results in Internet Appendix IA.E. Because $\Delta X_{i,t}$ is inherently more sensitive to timing and transitory noise than $X_{i,t}$, I view these estimates as supplementary. The appendix results indicate that emission changes can be associated with return predictability at short horizons, while the level effect is more stable across specifications and horizons.

Overall, the evidence rejects the restrictions implied by intensity sufficiency. Scale-invariance tests (Table 3) reject homogeneity of the intensity premium across firm size. Spanning tests (Tables 4–5) show that emission levels carry independent explanatory power conditional on firm value, with the rejection most pronounced among smaller firms and attenuating monotonically with scale. The rejection implies that any candidate characteristic for priced climate exposure must satisfy restrictions strictly stronger than those implied by intensity alone, and that these restrictions are empirically binding in the current sample. The framework developed in Sections 3 and 5 thus provides not only a diagnostic for existing measures but also a disciplining constraint on the construction of future ones: scalar characteristics are admissible only to the extent that they recover the underlying exposure function up to scale, and the present evidence indicates that emission intensity does not. Structurally estimating the components of $g(Z_{i,t})$ consistent with the observed cross-sectional patterns is a natural direction for future research.

7 Implications for Measurement, Allocation, and Policy

In this section, I develop implications of the exposure-based framework for measurement, capital allocation, and policy interpretation. The central insight is that asset prices reflect normalized exposure functions $g(Z_{i,t})/V_{i,t}$, while commonly used characteristics capture projections of this exposure onto selected metrics. When the chosen characteristic does not recover the underlying exposure function, this misalignment distorts both the measurement of climate risk and the allocation of capital. The subsections below illustrate this mechanism in a sequence of settings. Section 7.1 shows how firms' endogenous responses, such as efficiency improvements, can alter exposure in ways that are not captured by emission-based measures, leading to capital misallocation. Section 7.2 demonstrates that when multiple economic mechanisms jointly determine exposure, priced risk is inherently multidimensional, so that scalar measures, including emission intensity and composite ESG scores, recover at best projections of the underlying exposure functions. Finally, Section 7.3 revisits the policy benchmark to clarify why emission intensity arises naturally in implementation, even when it does not summarize equilibrium risk premia. Together, these results show that the choice of characteristic is not innocuous: it shapes both empirical inference and real economic outcomes.

7.1 Capital allocation and firm behavior

This section illustrates how misalignment between measured characteristics and underlying exposure functions can distort capital allocation and firm behavior. In particular, I show how endogenous energy efficiency can undermine emission intensity as a proxy for priced transition exposure. The goal is not to build a full production model, but to isolate a simple mechanism through which firms can improve commonly used greenness metrics while increasing their value-relevant exposure to priced climate or policy shocks.

The key insight is that firms respond to the characteristics that are measured and priced. When these characteristics do not coincide with the underlying exposure function $g(Z_{i,t})$, firms may optimally adjust behavior in ways that improve measured metrics but increase true exposure. A Jevons-type rebound provides a simple and economically relevant illustration of this mechanism. Improvements in energy efficiency lower marginal production costs, which induces firms to expand output. The expansion of production increases total energy use and emissions, potentially raising exposure to climate or policy shocks even as emissions per unit of output fall. As a result, firms

that appear greener according to intensity measures can face higher equilibrium risk premia by increasing their effective exposure to priced shocks.

While the section is illustrative, the mechanism has direct implications for capital allocation. If investors rely on intensity-based measures that do not capture the relevant exposure function, capital may be directed toward firms that improve measured metrics without reducing, and potentially increasing, underlying transition risk.

Efficiency choice and rebound. At date t, firm i uses energy $E_{i,t}$ and chooses an efficiency parameter $a_{i,t} > 0$ to produce an energy service

$$S_{i,t} = a_{i,t}E_{i,t}, \qquad X_{i,t} = \gamma E_{i,t},\ \gamma > 0,$$

so that $a_{i,t}$ increases service per unit of energy (and per unit of emissions). Let P_t^E denote the real energy price and suppose operating profits admit the reduced form

$$\Pi_{i,t} = \pi_i(S_{i,t}) - P_t^E E_{i,t} - \Psi_i(a_{i,t}),$$

with π_i increasing and concave and Ψ_i the private cost of efficiency. For given $(P_t^E, a_{i,t})$, optimal energy demand $E_{i,t}^*(a_{i,t})$ satisfies

$$a_{i,t}\pi_i'(a_{i,t}E_{i,t}^*) = P_t^E.$$

A Jevons-type rebound region is characterized by

$$\frac{\partial E_{i,t}^*}{\partial a_{i,t}} \geq 0,$$

so that higher efficiency does not reduce, and may increase, energy use and therefore emissions $X_{i,t} = \gamma E_{i,t}^*(a_{i,t})$. Because emissions are induced by the equilibrium energy choice, any shock sensitivity expressed in the primitives $(E_{i,t}^*, a_{i,t}, Z_{i,t})$ is equivalently a function of $(X_{i,t}, a_{i,t}, Z_{i,t})$, and its comparative statics with respect to efficiency need not be negative in the rebound region. Empirically, this implies that intensity can improve mechanically while priced exposure rises if efficiency shifts the firm into the rebound region, so intensity-based greenness need not be informative about the relevant exposure function $g(Z_{i,t})$.

Pricing with efficiency-dependent exposures. Let ζ_{t+1} denote an aggregate climate or policy shock and let $\varepsilon_{E,t+1}$ denote an energy-price (or energy-demand) shock. Write dividends as

$$D_{i,t+1} = \tilde{D}_{i,t+1} + g_\zeta(X_{i,t}, a_{i,t}, Z_{i,t})\,\zeta_{t+1} + g_E(X_{i,t}, a_{i,t}, Z_{i,t})\,\varepsilon_{E,t+1}, \tag{7.1}$$

where $Z_{i,t}$ collects other characteristics and $\tilde{D}_{i,t+1}$ is orthogonal (conditionally) to $(\zeta_{t+1}, \varepsilon_{E,t+1})$. The Jevons mechanism matters because efficiency affects energy demand and therefore both emissions and shock sensitivity, so the exposure coefficients in (7.1) generically depend on $(X_{i,t}, a_{i,t})$ jointly. More generally, efficiency can affect exposure not only through emissions but also through the economic incidence of climate shocks—for example through pass-through, adjustment costs, or regulatory thresholds—so that $g_\zeta(X_{i,t}, a_{i,t}, Z_{i,t})$ need not be proportional to emissions alone.

Under the same linear-factor pricing logic as in Proposition 4, expected excess returns load on value-normalized exposures:

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \lambda_t^{(\zeta)} \frac{g_\zeta(X_{i,t}, a_{i,t}, Z_{i,t})}{V_{i,t}} + \lambda_t^{(E)} \frac{g_E(X_{i,t}, a_{i,t}, Z_{i,t})}{V_{i,t}}, \tag{7.2}$$

with equilibrium prices of risk $(\lambda_t^{(\zeta)}, \lambda_t^{(E)})$ determined by preferences, risk sharing, and the aggregate exposure of cash flows and (where relevant) unspanned damages to $(\zeta_{t+1}, \varepsilon_{E,t+1})$. Equation (7.2) makes clear that the priced characteristics are the normalized exposure functions $g_\zeta(\cdot)/V_{i,t}$ and $g_E(\cdot)/V_{i,t}$, which depend on efficiency and emissions jointly. Once $a_{i,t}$ shifts energy use and hence shock sensitivity, there is no generic reason for premia to collapse to a function of $X_{i,t}/V_{i,t}$ alone.

A simple example with reversed comparative statics

To see how a rebound can raise premia, consider

$$\pi_i(S) = pS - \frac{c}{2}S^2, \qquad p > 0,\ c > 0,$$

and ignore $\Psi_i(a)$ for the moment. The optimal energy choice is

$$E_{i,t}^*(a_{i,t}) = \frac{a_{i,t}p - P_t^E}{ca_{i,t}^2} \quad \text{for } a_{i,t}p > P_t^E,$$

so that

$$\frac{\partial E_{i,t}^*}{\partial a_{i,t}} = \frac{2P_t^E - a_{i,t}p}{ca_{i,t}^3}.$$

Hence, for $a_{i,t} \leq 2P_t^E/p$, higher efficiency increases energy use and emissions.

Suppose the climate shock enters dividends through emission levels,

$$g_\zeta(X_{i,t}, a_{i,t}, Z_{i,t}) = \eta X_{i,t} = \eta\gamma E^*_{i,t}(a_{i,t}), \qquad g_E \equiv 0,$$

and the equilibrium price of climate risk satisfies $\lambda_t^{(\zeta)} > 0$. Then (7.2) reduces to

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \lambda_t^{(\zeta)} \frac{\eta\gamma E^*_{i,t}(a_{i,t})}{V_{i,t}}.$$

To isolate the exposure channel emphasized by the Jevons mechanism, consider the comparative static locally holding firm value $V_{i,t}$ fixed. This partial-equilibrium normalization captures how efficiency alters the firm's priced cash-flow exposure through equilibrium energy use and emissions, abstracting from the separate effect of efficiency on firm valuation through expected profits. In general equilibrium, changes in $a_{i,t}$ also affect $V_{i,t}$, but the exposure channel alone is sufficient to demonstrate that efficiency need not reduce priced climate exposure.

Under this normalization, in the rebound region,

$$\frac{\partial}{\partial a_{i,t}}\big(\mathbb{E}_t[R_{i,t+1}] - R_{f,t}\big)\Big|_{V_{i,t}} = \lambda_t^{(\zeta)} \frac{\eta\gamma}{V_{i,t}} \frac{\partial E^*_{i,t}}{\partial a_{i,t}} \geq 0,$$

with strict inequality for $a_{i,t} < 2P_t^E/p$. Thus, higher efficiency can increase equilibrium premia by raising emissions and hence normalized exposure to the priced climate factor. The full general-equilibrium comparative static additionally depends on how efficiency affects firm valuation through expected profits and continuation values, but the exposure channel alone is sufficient to demonstrate that efficiency improvements need not reduce priced climate exposure.

At the same time, higher efficiency increases service output per unit of energy and therefore reduces physical emissions intensity (emissions per unit of service), even as equilibrium emissions and priced exposure rise. As a result, physical or accounting intensity measures can improve while the priced exposure determining expected returns increases, so intensity-based rankings need not align with equilibrium risk premia in rebound regions.

The Jevons mechanism illustrates how improvements in measured characteristics, such as emission intensity, can coincide with increases in value-relevant exposure to priced climate risk. When firms adjust efficiency in ways that affect energy use and emissions, the resulting exposure functions depend jointly on emissions and technological choices, so there is no general reason for priced char-

acteristics to collapse to $X_{i,t}/V_{i,t}$ alone. This disconnect between measured metrics and underlying exposure is central to the misallocation channel discussed in this section.

7.2 Multiple exposure channels and the limits of aggregation

The analysis in Section 5 shows that different economic mechanisms map firm characteristics into distinct exposure functions. In practice, these mechanisms are likely to operate jointly, so that firms are exposed to multiple climate-related channels simultaneously. In such environments, expected returns take the form

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \sum_k \lambda_{k,t} \frac{g_k(Z_{i,t})}{V_{i,t}},$$

where each $g_k(\cdot)$ corresponds to a distinct channel through which climate or policy shocks affect cash flows.

When multiple exposure channels are present, the normalized exposure vector $\{g_k(Z_{i,t})/V_{i,t}\}_k$ is inherently multidimensional. There is therefore no reason for priced exposure to collapse to a scalar function of any single firm characteristic. Any univariate measure recovers at best a projection of this multidimensional object.

This observation has direct implications for the interpretation of ESG ratings and other composite sustainability measures. In practice, ESG scores aggregate a wide range of firm characteristics—including emissions, energy efficiency, governance indicators, and disclosure metrics—into a single index. Such aggregation implicitly assumes that these characteristics measure the same underlying exposure object, or at least that they are proportional to one another.

From the perspective of the exposure-based framework, however, different characteristics generally correspond to different underlying economic mechanisms and therefore to different exposure functions. Emission levels capture direct exposure to transition costs under carbon pricing, while energy efficiency, technology choices, or sector classifications reflect distinct channels through which firms are affected by climate-related shocks. When these mechanisms operate jointly, the corresponding exposure functions need not be proportional, and aggregation across them does not recover the relevant priced exposure.

This perspective also provides a structural interpretation for the pervasive divergence documented in ESG ratings. Berg et al. (2022) show that ratings from major providers correlate only weakly with one another, and they attribute the bulk of this divergence to differences in what providers choose to measure (*scope*) and how they weight individual indicators (*weight*), with

smaller contributions from measurement of common indicators. The exposure-based framework rationalizes this pattern: if different providers implicitly emphasize different underlying mechanisms—some weighting direct emissions exposure more heavily, others weighting technology or governance channels—their composite scores correspond to distinct projections of the multidimensional exposure vector $\{g_k(Z_{i,t})/V_{i,t}\}_k$ onto different subsets of firm characteristics. Divergence across raters is therefore not merely a symptom of measurement noise to be averaged away, but a direct consequence of the fact that no single scalar index can recover priced exposure when mechanisms operate jointly. Aggregating further across providers compounds rather than resolves the identification problem.

A related implication applies to efforts to recover a "true" greenium by combining multiple emission-based characteristics. Eskildsen et al. (2024) argue that standard carbon measures yield divergent estimates of the greenium because of measurement error and construct a composite green score by averaging across levels, intensities, and related characteristics. Under the exposure-based framework, such aggregation reduces measurement error only if the underlying signals correspond to the same priced object. When levels, intensities, and changes instead correspond to distinct exposure functions—as Section 5 shows arise naturally from per-unit policies, trajectory-based regulation, and net-zero benchmarks—averaging across them projects a multidimensional exposure object onto a scalar index and obscures rather than sharpens the economic interpretation of the resulting premium. The spanning and scale-invariance tests in Section 6.2 reinforce this point: emission levels and intensity are not noisy measurements of a common underlying object, but carry independent cross-sectional information about expected returns, a pattern that is inconsistent with the common-factor structure that motivates aggregation.

The issue can be illustrated using three simple mechanisms. Under a carbon tax, exposure is proportional to emission levels,

$$g(Z_{i,t}) = \eta X_{i,t}.$$

Under a scale-dependent policy, exposure may be nonlinear in emissions,

$$g(Z_{i,t}) = X_{i,t}^2.$$

Under a net-zero constraint, exposure may depend on deviations from a regulatory benchmark X^{NZ},

$$g(Z_{i,t}) = X_{i,t} - X^{NZ}.$$

Each mechanism implies a different mapping from firm characteristics into priced exposure. Aggregating across these measures—for example, by constructing a composite index that averages their normalized counterparts—does not correspond to any underlying exposure function $g(Z_{i,t})$. Instead, it combines characteristics associated with distinct economic mechanisms into a single index.

More generally, aggregation across characteristics associated with different mechanisms produces a projection of multidimensional exposure onto a scalar index. Unless the underlying exposure functions are proportional, this projection does not preserve the cross-sectional variation in priced exposure that determines equilibrium returns. As a result, variation in composite ESG scores may mix fundamentally different sources of risk and need not correspond to variation in priced climate exposure. The divergence documented by Berg et al. (2022) and the aggregation strategy of Eskildsen et al. (2024) are therefore two manifestations of the same underlying issue: in the absence of a theory of which characteristic corresponds to priced exposure, neither disagreement across raters nor averaging across indicators can be resolved by additional data alone.

Heterogeneity in investor preferences further amplifies this issue. As shown in Section 3.3, equilibrium prices of risk reflect a weighted aggregation of heterogeneous valuations across distinct exposure channels. In a multi-channel environment, both the mapping from firm characteristics into priced exposure and the pricing of that exposure depend on the interaction between multiple mechanisms and heterogeneous preferences.

These considerations reinforce the central implication of the paper: identifying priced climate risk requires recovering the underlying exposure functions rather than relying on scalar proxies that aggregate across heterogeneous channels.

7.3 Planner, social cost of carbon, and a Pigouvian tax

This subsection provides a policy benchmark that clarifies when emission intensity can serve as a valid summary of climate exposure. In the complete-markets environment underlying sustainable-discount-rate approaches, the relevant object is a deterministic implementation wedge: policy-induced distortions to firms' cash flows scale linearly with emissions and are expressed on a per-unit-of-value basis. As a result, normalized exposure collapses to emission intensity. In this setting, intensity is not a projection but the correct underlying object, and therefore emerges naturally as a sufficient statistic for implementing policy objectives. This benchmark highlights that the role of emission intensity depends on the structure of the exposure function. In the equilibrium

asset-pricing environment studied in the main analysis, exposure functions $g(Z_{i,t})$ are determined by technology, policy design, and firm behavior and need not be proportional to emissions. As a result, intensity typically reflects a projection of multidimensional exposure rather than the exposure itself. The planner benchmark therefore clarifies why intensity-based measures are prevalent in policy design, while also making clear that their validity hinges on restrictive conditions that need not hold in equilibrium pricing.

Formally, consider again dates $t = 0, 1$ and a single consumption good. A representative agent has utility

$$U = u(C_0) + \beta\Big(u(C_1) - d(X_1)\Big), \tag{7.3}$$

where C_0 and C_1 denote aggregate consumption at dates 0 and 1, $\beta \in (0, 1)$ is the subjective discount factor, and $d : \mathbb{R}_+ \to \mathbb{R}_+$ is a strictly increasing, convex damage function of aggregate emissions $X_1 \geq 0$. The function $d(\cdot)$ can be interpreted as the expected disutility from climate damages, for example $d(X_1) = \mathbb{E}[\Gamma(\zeta) \mid X_1]$ in a reduced-form sense. The per-period utility u is strictly increasing and concave.

Firm i produces a deterministic date-1 payoff $Y_i \geq 0$ per unit of scale and emits $X_i \geq 0$ per unit of scale. If I_i denotes the scale of firm i at date 0, we can aggregate emissions as

$$X_1 = \sum_{i=1}^{N} X_i I_i. \tag{7.4}$$

Let G_0 denote the date-0 endowment of the consumption good. The resource constraints are

$$C_0 + \sum_{i=1}^{N} I_i \leq G_0, \qquad C_1 = \sum_{i=1}^{N} Y_i I_i. \tag{7.5}$$

The technological mapping from I_i to (Y_i, X_i) is kept implicit. (Y_i, X_i) are treated as primitive characteristics of one unit of firm i, and the planner chooses the scale I_i.

The social planner chooses $(C_0, C_1, \{I_i\}_{i=1}^N)$ to maximize (7.3) subject to (7.4)-(7.5). Let λ_0 and λ_1 denote the Lagrange multipliers on the two resource constraints. The planner's Lagrangian is

$$\mathcal{L} = u(C_0) + \beta\big(u(C_1) - d(X_1)\big) + \lambda_0\Big(G_0 - C_0 - \sum_i I_i\Big) + \lambda_1\Big(\sum_i Y_i I_i - C_1\Big).$$

The first-order conditions for C_0 and C_1 are

$$u'(C_0) = \lambda_0, \qquad \beta u'(C_1) = \lambda_1. \tag{7.6}$$

The first-order condition for I_i is

$$-\lambda_0 + \lambda_1 Y_i - \beta d'(X_1) X_i = 0. \tag{7.7}$$

Combining (7.6) and (7.7) yields

$$\frac{\lambda_0}{\lambda_1} = Y_i - \frac{\beta d'(X_1)}{\lambda_1} X_i = Y_i - \frac{d'(X_1)}{u'(C_1)} X_i. \tag{7.8}$$

The shadow cost of an additional unit of emissions in terms of the consumption numeraire is defined as the social cost of carbon,

$$S \equiv \frac{d'(X_1)}{u'(C_1)}. \tag{7.9}$$

Expression (7.8) can then be written as

$$\frac{\lambda_0}{\lambda_1} = Y_i - S X_i. \tag{7.10}$$

Equation (7.10) shows that, at the planner optimum, a marginal increase in scale I_i raises date-1 resources by Y_i but also increases the planner's marginal damage by SX_i. It is therefore convenient to interpret $Y_i - SX_i$ as the planner-implied net date-1 payoff (in consumption units) generated by one unit of firm i's scale.

To express this condition in present-value form, define the planner-implied intertemporal discount factor as

$$q \equiv \frac{\lambda_1}{\lambda_0} = \beta \frac{u'(C_1)}{u'(C_0)}. \tag{7.11}$$

Then (7.10) is equivalent to

$$1 = q(Y_i - S X_i). \tag{7.12}$$

Thus, the date-0 shadow value of one unit of scale is $q(Y_i - SX_i)$, the present value of its net date-1 payoff.

For later comparison, it is convenient to parameterize the planner-implied discount factor by a

(gross) risk-free rate r via

$$q \equiv \frac{\lambda_1}{\lambda_0} = \frac{1}{1+r}. \tag{7.13}$$

A Pigouvian carbon tax τ per unit of emissions, levied uniformly on firms and rebated lump-sum to the representative household, decentralizes the planner's allocation if $\tau = S$. Under such a tax, firm i's date-1 net payoff per unit of scale is $Y_i - \tau X_i$, and its date-0 value is

$$V_i = q(Y_i - \tau X_i) = \frac{Y_i - \tau X_i}{1+r}. \tag{7.14}$$

If $\tau = S$, then (7.12) and (7.14) coincide, implying that the competitive equilibrium with tax $\tau = S$ reproduces the first-best allocation.

7.3.1 Sustainable discount rates and the emission-intensity wedge

The sustainable-discount-rate approach asks whether the same allocation that would be achieved by a Pigouvian carbon tax can instead be implemented by adjusting firms' discount rates while leaving their cash flows untaxed. In the one-period environment above, this equivalence can be made explicit and directly related to emission intensity.

Fix a firm i and suppose that, under the carbon tax τ, its date-0 value is V_i given by (7.14). Let r denote the financial discount rate that applies in the absence of any climate-related discount-rate adjustment. A sustainable discount rate r_i^x for firm i is defined implicitly by the condition that the present value of the taxed cash flow discounted at r equals the present value of the untaxed cash flow discounted at r_i^x, that is,

$$\frac{Y_i - \tau X_i}{1+r} = \frac{Y_i}{1+r_i^x}. \tag{7.15}$$

In this deterministic two-date environment, (7.15) admits a closed-form solution that links r_i^x to emission intensity exactly. From (7.14), $V_i(1+r) = Y_i - \tau X_i$, so $Y_i = V_i(1+r) + \tau X_i$. Substituting into (7.15) and using $V_i = (Y_i - \tau X_i)/(1+r)$ gives

$$1 + r_i^x = \frac{Y_i}{V_i} = \frac{V_i(1+r) + \tau X_i}{V_i} = (1+r) + \tau \frac{X_i}{V_i},$$

so that

$$r_i^x - r = \tau \frac{X_i}{V_i}. \tag{7.16}$$

The wedge between the sustainable and financial discount rates is therefore proportional to the

firm's emission intensity, with the coefficient given by the carbon tax τ. The relation in (7.16) is an exact identity in this benchmark, with V_i denoting the firm's value under the tax as in (7.14).

If the tax τ coincides with the SCC S in (7.9), then (7.16) links the sustainable discount rate directly to the SCC and emission intensity:

$$r_i^x = r + S\frac{X_i}{V_i}. \tag{7.17}$$

More generally, suppose that the prevailing carbon tax is τ while the planner's SCC is S, and consider the additional discount-rate wedge required to move firm value from its level under tax τ to its level under the first-best tax S. Equating the change in firm value induced by a tax change $\mathrm{d}\tau$ to that induced by a discount-rate change $\mathrm{d}r_i$,

$$\mathrm{d}V_i^{(\tau)} = -\frac{X_i}{1+r}\mathrm{d}\tau, \qquad \mathrm{d}V_i^{(r)} = -\frac{V_i}{1+r}\mathrm{d}r_i,$$

and setting $\mathrm{d}V_i^{(\tau)} = \mathrm{d}V_i^{(r)}$ yields $\mathrm{d}r_i = (X_i/V_i)\,\mathrm{d}\tau$. Applying this with $\mathrm{d}\tau = S - \tau$ and linearizing around the initial allocation gives

$$r_i^x \approx r + (S-\tau)\frac{X_i}{V_i}. \tag{7.18}$$

The approximation in (7.18) is local in the size of the policy gap $S - \tau$; when $\tau = 0$, the relation reduces to the exact identity (7.17) with the wedge given by the full SCC.

Equations (7.9) and (7.18) formalize the sustainable-discount-rate representation developed by Pedersen (2025). The SCC S is the planner's shadow value of emissions in units of the consumption numeraire, and the same allocation that would be achieved by adjusting the carbon tax from τ to S can be implemented by tilting each firm's discount rate in proportion to its emission intensity X_i/V_i. In this complete-markets benchmark, X_i/V_i is the sufficient statistic for how strongly a firm's cost of capital must be adjusted to implement the planner's allocation.

This is an implementation result rather than a restriction on equilibrium pricing. In the language of Proposition 3, the planner environment corresponds to the linear-sufficiency case: the per-ton policy instrument implies that homogeneity, scale invariance, and linearity all hold by construction of the implementation wedge, rather than as restrictions on equilibrium cash-flow exposure. Equation (7.18) therefore characterizes the deterministic discount-rate wedge required to decentralize a first-best allocation under complete markets. The appearance of emission intensity is a normalization, not a pricing statement. In contrast, equilibrium asset pricing depends on how cli-

mate and policy shocks enter firms' cash flows through exposure functions $g(Z_{i,t})$. Once one allows for heterogeneous investors and incomplete risk sharing, the characteristics that govern equilibrium risk premia need not collapse to X_i/V_i and can instead depend on emission levels and more general functions of firm characteristics. The planner benchmark therefore clarifies why emission intensity emerges naturally in policy implementation, while highlighting that equilibrium risk premia depend on underlying exposure functions rather than on intensity alone.

8 Conclusion

In this paper, I develop an equilibrium model of climate risk premia. Climate risk premia are determined by two objects: the equilibrium price of risk, λ, and the economic mechanism that maps firm characteristics into cash-flow exposure, $g(Z_{i,t})$. The price of risk reflects aggregate conditions, such as policy uncertainty, risk aversion, total market emissions, and investor composition, while the exposure function is shaped by firm technology, policy design, and firm behavior. Under incomplete markets, the price of risk can be decomposed into components reflecting compensation for climate risk and hedging demand, but does not determine which firm characteristics govern exposure. In turn, the form of the exposure function is independent of market completeness.

This separation between the price of risk and the exposure function implies that identifying climate risk premia requires specifying the relevant exposure function. I characterize the conditions under which a given characteristic summarizes climate risk premia using emission intensity as an example. Emission intensity is sufficient only when the exposure function satisfies homogeneity and scale-invariance conditions. These conditions arise mechanically in planner implementations, where wedges are proportional to emissions per unit of firm value, and reflect restrictions on the exposure function rather than features of market completeness, but do not generally hold in equilibrium asset pricing. More generally, commonly used carbon metrics need not coincide with the normalized exposure $g(Z_{i,t})/V_{i,t}$ implied by the underlying mechanism. In this sense, the choice of characteristic is not a measurement decision but a statement about the underlying economic mechanism.

This distinction clarifies how to select empirically relevant characteristics. A valid characteristic must correspond to the normalized exposure implied by the underlying mechanism. Absent this mapping, regression coefficients reflect projections of equilibrium risk premia onto the chosen metric rather than structural parameters. Empirical tests reject the sufficiency restrictions, suggesting that commonly used metrics do not fully capture the relevant exposure functions and that part of the

variation in estimates across the literature may reflect differences in the underlying economic objects being measured, not sampling variation or model selection.

A related methodological point concerns what statistical significance diagnoses in this setting. A regression of returns on any candidate characteristic recovers the partial projection of equilibrium climate risk premia onto that characteristic, conditional on controls. This projection is generically non-zero whenever the characteristic carries information about the true exposure, so a significant coefficient is consistent with the characteristic being structurally correct, with it being a projection of a different correctly specified object, or with it being correlated with the truth through a nuisance channel. An insignificant coefficient is similarly ambiguous: it can reflect a characteristic that is uninformative about the exposure function, but it can also reflect a small equilibrium price of climate risk in-sample, inadequate power, or a mechanism that is latent rather than active. Standard significance tests therefore cannot, on their own, adjudicate which characteristic recovers the relevant exposure. The structural restrictions derived in Proposition 3 are what discipline the mapping from coefficients to exposure, because they test implications that do not depend on the level of the risk premium.

An important implication of my framework is that misidentifying the relevant exposure function leads to distorted pricing, capital allocation, and firm incentives. When investors or regulators rely on characteristics that do not correspond to the underlying exposure, equilibrium prices reflect projections onto the chosen metric rather than true risk. This mispricing directs capital toward firms that appear low-risk under the selected characteristic, even if their underlying exposure remains high, and induces firms to optimize the measured metric rather than the underlying exposure. As a result, decarbonization efforts may improve reported indicators without reducing transition risk. This logic extends directly to sustainable finance practice: regulatory taxonomies and disclosure standards often emphasize particular metrics, implicitly assuming the conditions under which those measures are sufficient for pricing. When these conditions do not hold, portfolios that satisfy regulatory or disclosure-based criteria can nonetheless carry substantial climate transition risk if the underlying exposure depends on dimensions the metric does not capture.

Beyond metric choice, the framework provides a micro-foundation for the green and brown classifications central to the climate-finance literature (e.g., Pastor et al. (2021)). Rather than treating these types as exogenous labels, the analysis endogenizes them as equilibrium exposure functions shaped by technology and fundamentals. A firm's climate-risk profile is therefore a dynamic object: changes in efficiency, production structure, or policy exposure can alter priced risk

even when conventional ESG metrics improve.

The framework provides a disciplined approach to empirical work. Emission-level specifications are not arbitrary alternatives to intensity but arise from specific exposure functions implied by the underlying mechanism. More generally, testing scale-invariance and spanning restrictions offers a template for evaluating whether a proposed climate-risk metric is theoretically coherent. When these restrictions fail, recovering the underlying exposure functions becomes necessary for interpreting empirical premia. Future research should therefore structurally estimate the exposure functions $g_k(Z_{i,t})$ implied by specific policy and technology channels, quantify the relative importance of transition risk and preference-driven greenium components, and extend the framework to incorporate dynamic investment and abatement decisions.

References

Berg, F., J. F. Kolbel, and R. Rigobon (2022). Aggregate confusion: The divergence of esg ratings. *Review of Finance 33*(2), 1315–1344.

Bolton, P. and M. Kacperczyk (2021). Do investors care about carbon risk? *Journal of Financial Economics 142*(2), 517–549.

Bolton, P. and M. Kacperczyk (2023). Global pricing of carbon-transition risk. *The Journal of Finance 78*(6), 3677–3754.

Cenedese, G., S. Han, and M. Kacperczyk (2024). Carbon transition risk and net-zero portfolios. *Working Paper*.

Chittaro, L., M. Piazzesi, M. Sena, and M. Schneider (2025). Asset returns as carbon taxes. *Working Paper*.

Dangl, T., M. Halling, J. Yu, and J. Zechner (2025). Social preferences and corporate investment. *Journal of Financial Economics*, forthcoming.

De Angelis, T., P. Tankov, and D. O. Zerbib (2023). Climate impact investing. *Management Science 69*(12), 7669–7692.

Eskildsen, M., I. Marcus, T. Jensen, and L. Pedersen (2024). In search of true greenium. *Working Paper*.

Hong, H. and M. Kacperczyk (2009). The price of sin: The effects of social norms on markets. *Journal of Financial Economics 93*(1), 15–36.

Hsu, P.-H., K. Li, and C.-Y. Tsou (2023). The pollution premium. *Journal of Finance 78*(3), 1343–1392.

Pastor, L., R. F. Stambaugh, and L. Taylor (2021). Sustainable investing in equilibrium. *Journal of Financial Economics 142*(2), 550–571.

Pastor, L., R. F. Stambaugh, and L. Taylor (2022). Dissecting green returns. *Journal of Financial Economics 142*(2), 550–571.

Pastor, L., R. F. Stambaugh, and L. Taylor (2025a). Carbon burden. *Working Paper*.

Pastor, L., R. F. Stambaugh, and L. Taylor (2025b). Sustainable investing. *Annual Review of Financial Economics 17*, 321–341.

Pedersen, L. (2025). Carbon pricing versus green finance. *The Journal of Finance*, forthcoming.

Pedersen, L., S. Fitzgibbons, and L. Pomorski (2021). Responsible investing: The esg-efficient frontier. *Journal of Financial Economics 142*(2), 572–597.

Sautner, Z., L. Van Lent, G. Vilkov, and R. Zhang (2023). Pricing climate change exposure. *Management Science 69*(2), 7540–7561.

Zerbib, D. O. (2022). A sustainable capital asset pricing model (s-capm): Evidence from environmental integration and sin stock exclusion. *Review of Finance 26*(6), 1345–1388.

Table 1: Using the model: priced objects under alternative exposure mechanisms

Economic mechanism	**Exposure structure** $g(Z_{i,t})$	**Object priced in equilibrium**	**Regression object**
Per-unit transition risk (taxes, permits, linear policy shocks)	$g(Z_{i,t}) = \eta X_{i,t}$	Emission scale, conditional on firm value	$X_{i,t}$
Planner implementation (complete markets, SCC-based policy)	Deterministic loss $S_t X_{i,t}$	Firm-specific discount-rate wedge	$X_{i,t}/V_{i,t}$
Scale-invariant equilibrium (homogeneous exposure benchmark)	$g(Z_{i,t}) = \kappa_t X_{i,t}$	Priced exposure identical for firms with equal intensity	$X_{i,t}/V_{i,t}$
Trajectory-based regulation (policy responds to changes or trends)	$g(Z_{i,t}) = \phi_1 X_{i,t} + \phi_2 \Delta X_{i,t}$	Emission levels and changes	$(X_{i,t},\, \Delta X_{i,t})$
Path-based constraints / net-zero (benchmarks, phase-out schedules)	$g(Z_{i,t}) = \vartheta_1 X_{i,t} + \vartheta_2 (X_{i,t} - X_{i,t}^{NZ})$	Emissions and distance to targets	$(X_{i,t},\, X_{i,t} - X_{i,t}^{NZ})$

Table 2: **Summary statistics**

The sample includes listed companies from 92 countries over 2005-2022. *Returns* are monthly stock returns (in percent). Scope 1 emission levels (X) are in metric tons of CO_2 equivalent. Market value of equity (V) is in millions of US dollars. Emission intensity (X/V) is the ratio of emissions to market value of equity, winsorized at the 2.5% level. Scope 1 growth rate is the year-over-year percentage change in emissions. Statistics reported are mean, median, and standard deviation across all observations.

Variable / Statistic	Mean	Median	St. dev.
Returns (in %)	0.890	0.348	12.134
Scope 1 levels (X)	1,696,164	26,141	11,700,000
Market value of equity (V)	735,137	12,221	8,774,292
Intensity (X/V)	0.636	0.018	1.967
Log(X)	10.356	10.174	2.850
Log(V)	9.623	9.411	2.463
Log(X/V)	0.732	0.616	3.138
Scope 1 growth rate	0.199	-0.041	2.869
Number of firms: 13,373			
Number of countries: 92			
Number of industries: 438			

Table 3: **Predicting Stock Returns with Emission Intensity and Tests of Scale Invariance**

The sample is at the firm-month level and includes listed companies from 92 countries over 2005-2022. The dependent variable is monthly stock returns ($Returns_{i,t+1}$). Each period, firms are sorted by market value of equity into quintiles ($Size\ bin$). Emission intensity (X/V) is the ratio of Scope 1 emissions to market value. All regressions include industry × time fixed effects and country fixed effects. Standard errors (in parentheses) are double-clustered by industry and time. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. The bottom row reports p-values from Wald tests of the null hypothesis that coefficients on intensity are equal across all five size bins.

	(1)	(2)
VARIABLES	Returns	
Log(X/V)	0.208^{***}	0.176^{***}
	(0.056)	(0.036)
Size bin 2	0.039	0.100
	(0.176)	(0.108)
Size bin 3	−0.178	−0.109
	(0.217)	(0.116)
Size bin 4	−0.257	−0.176
	(0.201)	(0.142)
Size bin 5	-0.514^{**}	-0.377^{**}
	(0.213)	(0.178)
Size bin 2 × Log(X/V)	-0.049^{*}	-0.046^{*}
	(0.026)	(0.025)
Size bin 3 × Log(X/V)	−0.045	−0.036
	(0.034)	(0.031)
Size bin 4 × Log(X/V)	-0.072^{*}	-0.064^{*}
	(0.034)	(0.031)
Size bin 5 × Log(X/V)	−0.045	−0.038
	(0.029)	(0.031)
Log(M/B)		-0.238^{***}
		(0.077)
Book leverage		−0.148
		(0.240)
Momentum		0.458^{***}
		(0.157)
Investment/Assets		0.182
		(0.835)
Log(PPE)		−0.007
		(0.016)
Volatility		2.169
		(3.032)
ROE		0.753^{**}
		(0.267)
Age		0.602
		(0.970)
Constant	0.924^{***}	0.783^{**}
	(0.130)	(0.285)
Industry × Time FE	Yes	Yes
Country FE	Yes	Yes
Observations	1,004,760	1,004,760
R-squared	0.236	0.236
Wald test p-value	0.049	0.023

Table 4: **Predicting Stock Returns with Emission Levels, Market Values, and Scale**

The sample is at the firm-month level and includes listed companies from 92 countries over 2005-2022. The dependent variable is monthly stock returns ($Returns_{i,t+1}$) over the next month (columns 1-3) or next 12 months (columns 4-6). Emission levels (X) is Scope 1 emissions; V is the firm market capitalization. All regressions include industry × time fixed effects and country fixed effects. Standard errors (in parentheses) are double-clustered by industry and time. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. The bottom row reports the p-values of the following Wald tests. Columns 1 and 4 test the restriction $\beta_{logX} + \beta_{logV} = 0$. Columns 2 and 5 test $\beta_{logX} + \beta_{logSales} = 0$. Columns 3 and 6 jointly test both restrictions. All tests are Wald tests based on the clustered covariance matrix.

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	Returns					
	1 month ahead			12 months ahead		
Log(X)	0.125***	0.077***	0.111***	0.053***	0.009	0.031**
	(0.020)	(0.018)	(0.019)	(0.012)	(0.013)	(0.011)
Log(V)	−0.270***		−0.291***	−0.143***		−0.174***
	(0.054)		(0.061)	(0.038)		(0.045)
Log(Sales)		−0.014	0.067**		0.056**	0.105***
		(0.022)	(0.030)		(0.023)	(0.030)
Log(M/B)	−0.192**	−0.388***	−0.177**	−0.122	−0.229**	−0.103
	(0.066)	(0.097)	(0.063)	(0.081)	(0.098)	(0.078)
Book leverage	−0.224	0.148	−0.235	−0.213	−0.007	−0.221
	(0.220)	(0.270)	(0.220)	(0.175)	(0.199)	(0.176)
Momentum	0.469***	0.432**	0.472***	−0.313	−0.327*	−0.304
	(0.157)	(0.160)	(0.155)	(0.179)	(0.182)	(0.180)
Investment/Assets	−0.210	0.965	0.075	0.173	1.175*	0.641
	(0.888)	(0.863)	(0.865)	(0.710)	(0.641)	(0.661)
Log(PPE)	0.043	−0.081***	0.015	0.037	−0.065***	−0.008
	(0.028)	(0.016)	(0.020)	(0.022)	(0.019)	(0.019)
Volatility	2.004	2.549	2.022	−0.495	−0.144	−0.481
	(3.123)	(3.147)	(3.119)	(1.381)	(1.363)	(1.379)
ROE	0.808**	0.738**	0.744**	0.647*	0.562	0.556
	(0.282)	(0.269)	(0.274)	(0.338)	(0.332)	(0.332)
Age	0.598	0.749	0.517	−0.106	−0.056	−0.181
	(0.932)	(0.906)	(0.912)	(0.930)	(0.912)	(0.938)
Constant	1.699***	0.683	1.630***	1.454***	0.782***	1.339***
	(0.454)	(0.398)	(0.445)	(0.318)	(0.193)	(0.297)
Observations	1,004,760	1,003,810	1,003,810	850,722	850,085	850,085
Industry×Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.236	0.236	0.237	0.257	0.257	0.257
Wald test p-value (single sum test)	0.008	0.009		0.031	0.006	
Wald test p-value (joint sum test)			0.001			0.001

Table 5: **Within-size-bin Spanning Tests**

The sample is at the firm-month level and includes listed companies from 92 countries over 2005-2022. The dependent variable is monthly stock returns ($Returns_{i,t+1}$). The table reports spanning regressions estimated separately within firm size quintiles, where size is measured by market value of equity. Each column estimates the regression in Table 4 using firms in the corresponding size bin. All regressions include the same set of controls as those in Table 4 as well as industry × time fixed effects and country fixed effects. Standard errors (in parentheses) are double-clustered by industry and time. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. The bottom row reports p-values from Wald tests of the homogeneity restriction $\beta_{logX} + \beta_{logV} = 0$ estimated within each size quintile. Size quintiles are formed each period using market value of equity.

	Firm size quintile (market equity)				
	(1) Smallest	(2)	(3)	(4)	(5) Largest
Log(X)	0.164**	0.158***	0.100***	0.137***	0.126***
	(0.065)	(0.022)	(0.021)	(0.028)	(0.019)
Log(V)	−0.380***	−0.428**	−0.221***	−0.270**	−0.195**
	(0.091)	(0.149)	(0.072)	(0.099)	(0.080)
Log(M/B)	−0.242*	0.016	−0.160	−0.107	−0.265***
	(0.120)	(0.092)	(0.120)	(0.063)	(0.083)
Book leverage	−0.081	−0.371	−0.181	−0.371	−0.308
	(0.284)	(0.281)	(0.370)	(0.232)	(0.273)
Momentum	0.882***	0.356	0.750***	0.580***	0.420*
	(0.260)	(0.210)	(0.249)	(0.163)	(0.203)
Investment/Assets	1.188	−0.771	1.076	−0.321	0.204
	(1.712)	(1.036)	(0.850)	(1.718)	(0.930)
Log(PPE)	0.032	0.023	0.048	0.070	−0.010
	(0.034)	(0.034)	(0.034)	(0.051)	(0.035)
Volatility	3.990	2.395	−0.789	−0.782	−0.927
	(4.176)	(3.994)	(2.047)	(1.995)	(1.871)
ROE	0.812*	0.460	0.851*	0.843***	0.684*
	(0.399)	(0.323)	(0.438)	(0.276)	(0.332)
Age	1.127	1.499	−1.583	−1.731	−0.441
	(1.429)	(1.174)	(1.634)	(1.167)	(2.312)
Constant	1.453**	2.654**	1.433*	1.597	2.120***
	(0.646)	(1.227)	(0.748)	(1.008)	(0.531)
Industry × Time FE	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes
Observations	180,863	184,655	188,139	187,180	189,729
R^2	0.346	0.356	0.338	0.324	0.339
Wald test p-value	0.004	0.075	0.089	0.238	0.326

Appendix

A.1 Proofs and derivations

Proof of Proposition 1. By (3.3), for any h that holds a positive amount of stock i we have

$$\frac{\mu_i - V_i R_f}{\eta X_i} = \alpha_h \sigma_\zeta^2 \delta_h.$$

Since this ratio must be equalized across all stocks i and all households h with interior holdings, it must equal a common scalar λ:

$$\frac{\mu_i - V_i R_f}{\eta X_i} = \lambda \quad \forall i,$$

which yields (3.4). Substituting back gives $\alpha_h \sigma_\zeta^2 \delta_h = \lambda$, which yields (3.5). □

Proof of Proposition 2. In the static CARA-normal environment, aggregate consumption net of damages is $C_{h,1} - \phi_h \Gamma_1 = C_{h,1} - \phi_h \kappa \bar{X}\, \zeta_1$, so each household's SDF is exponential-affine in the climate shock,

$$m_{h,1} = \beta \frac{u_c^h(C_{h,1} - \phi_h \Gamma_1)}{u_c^h(C_{h,0} - \phi_h \Gamma_0)} \propto \exp\big(\alpha_h \phi_h \kappa \bar{X}\, \zeta_1\big),$$

where the proportionality absorbs β and all terms measurable at date 0. With mean-variance/CARA aggregation, the equilibrium (aggregate) SDF that prices all traded assets is the risk-tolerance-weighted average of individual SDF loadings, hence

$$m_1 \propto \exp\big(B\, \zeta_1\big), \qquad B = \kappa \bar{X} \frac{\sum_{h=1}^H \frac{\phi_h}{\alpha_h}}{\sum_{h=1}^H \frac{1}{\alpha_h}}.$$

Firm i's payoff (dividend) loads on the same shock with exposure ηX_i in the baseline specification, so $D_i = \tilde{D}_i + \eta X_i \zeta_1$ with $\tilde{D}_i \perp \zeta_1$. Using the Euler equation $1 = \mathbb{E}[m_1 R_i]$ and a second-order (mean-variance) expansion yields the standard linear pricing relation

$$\mathbb{E}[R_i] - R_f = -\operatorname{Cov}(R_i, \log m_1) = \sigma_\zeta^2 \,(\eta X_i)\,\Big(1 - B\Big),$$

where we used $\log m_1 = \text{const} + B\zeta_1$ and $\operatorname{Cov}(\zeta_1, \zeta_1) = \sigma_\zeta^2$. Aggregating across firms gives the climate factor price

$$\lambda = \frac{\sigma_\zeta^2 \bar{X}}{\sum_{h=1}^H \frac{1}{\alpha_h}} \Big(\eta - \kappa \sum_{h=1}^H \phi_h\Big),$$

which is (3.8). □

Proof of Proposition 3. Condition (iii) implies that, for any fixed intensity level $u \in \mathbb{R}_+$, all firms with $X_{i,t}/V_{i,t} = u$ share the same exposure $g(Z_{i,t})/V_{i,t}$. Define $\varphi_t(u)$ to be this common value:

$$\varphi_t(u) \equiv \frac{g(Z_{i,t})}{V_{i,t}} \quad \text{for any firm } i \text{ with } \frac{X_{i,t}}{V_{i,t}} = u.$$

Condition (i) ensures that intensities span an open subset of $\mathbb{R}_+$, so that φ_t can be defined on an interval rather than on a finite set. This establishes (3.15).

To obtain (3.16), fix a firm with characteristics $(X_{i,t}, V_{i,t})$ and intensity $u = X_{i,t}/V_{i,t}$. By homogeneity of degree one in $(X_{i,t}, V_{i,t})$, for any $\lambda > 0$,

$$g(Z_{i,t}; \lambda X_{i,t}, \lambda V_{i,t}) = \lambda\, g(Z_{i,t}; X_{i,t}, V_{i,t}).$$

Dividing both sides by $\lambda V_{i,t}$ yields

$$\frac{g(Z_{i,t}; \lambda X_{i,t}, \lambda V_{i,t})}{\lambda V_{i,t}} = \frac{g(Z_{i,t}; X_{i,t}, V_{i,t})}{V_{i,t}}.$$

The left-hand side is the value-normalized exposure of a firm whose intensity remains u (since $(\lambda X_{i,t})/(\lambda V_{i,t}) = u$). By condition (iii), it depends only on u and therefore equals $\varphi_t(u)$, while the right-hand side equals $\varphi_t(u)$ by definition. Hence

$$\frac{g(Z_{i,t})}{V_{i,t}} = \varphi_t\left(\frac{X_{i,t}}{V_{i,t}}\right),$$

which restates (3.15).

Under the additional assumption in the proposition that the mapping is linear in size (equivalently, that $g(Z_{i,t})/V_{i,t}$ is proportional to $X_{i,t}/V_{i,t}$ on the relevant support), there exists a scalar κ_t such that for all firms,

$$\frac{g(Z_{i,t})}{V_{i,t}} = \kappa_t \frac{X_{i,t}}{V_{i,t}}.$$

Multiplying both sides by $V_{i,t}$ yields $g(Z_{i,t}) = \kappa_t X_{i,t}$ and establishes (3.16). Uniqueness in the sense of Definition 3.1 follows because any one-factor representation observationally equivalent to (3.14) must use a characteristic proportional to the priced characteristic $g(Z_{i,t})/V_{i,t}$, hence proportional to $X_{i,t}/V_{i,t}$ under (3.16).

Necessity. Conversely, suppose $g(Z_{i,t}) = \kappa_t X_{i,t}$ for all i. Condition (ii) holds since

$$g(Z_{i,t};\, \lambda X_{i,t},\, \lambda V_{i,t}) = \kappa_t(\lambda X_{i,t}) = \lambda\, \kappa_t X_{i,t} = \lambda\, g(Z_{i,t};\, X_{i,t},\, V_{i,t})\,.$$

Condition (iii) holds since $X_{i,t}/V_{i,t} = X_{j,t}/V_{j,t}$ implies

$$\frac{g(Z_{i,t})}{V_{i,t}} = \kappa_t \frac{X_{i,t}}{V_{i,t}} = \kappa_t \frac{X_{j,t}}{V_{j,t}} = \frac{g(Z_{j,t})}{V_{j,t}}.$$

Condition (i) is a nondegeneracy requirement on the cross-sectional support and is maintained throughout. More generally, if (3.15) holds for some measurable ϕ_t, then (iii) is immediate, and homogeneity of degree one (ii) holds whenever ϕ_t arises from a homogeneous underlying exposure function, which is the case in any representation that respects the dimensional structure of dollar cash-flow exposures. Conditions (i)–(iii), together with continuity and linearity of ϕ_t, are therefore jointly necessary and sufficient for emission intensity to be a unique sufficient statistic in the sense of Definition 3.1. □

□

A Internet Appendix: Robustness to nonlinear damages and multiple climate shocks

This section shows that the main asset-pricing implications of our framework are robust to nonlinear damage functions and to the presence of multiple climate shocks (for example, physical, policy, and technology shocks). In particular, I show that under standard small-shock approximations the pricing kernel remains affine in the climate factors, expected excess returns are proportional to X_i/V_i with a combined price of risk $\lambda^{(1)}\eta_1 + \lambda^{(2)}\eta_2$. In this case, emission levels determine the underlying dollar exposure, so that normalized exposure is proportional to X_i/V_i.

A.1 Nonlinear damages in a single climate factor

I first extend the static damage specification to allow for nonlinear dependence on the climate shock. For expositional clarity, I work in a one-period environment. The dynamic case follows by applying the same local arguments period by period.

Let ζ be a scalar climate factor with $\mathbb{E}[\zeta] = 0$ and $\mathrm{Var}(\zeta) = \sigma_\zeta^2 \in (0, \infty)$, and let aggregate emissions be $\bar{X} = \sum_i X_i$ as before. Instead of the linear damage function $\Gamma = \kappa \bar{X}\zeta$, suppose that damages are given by

$$\Gamma = h(\bar{X}, \zeta), \tag{A.1}$$

where $h : \mathbb{R}_+ \times \mathbb{R} \to \mathbb{R}$ is twice continuously differentiable and satisfies $h(\bar{X}, 0) = 0$. A quadratic specification, for example,

$$h(\bar{X}, \zeta) = \kappa_1 \bar{X}\zeta + \frac{1}{2}\kappa_2 \bar{X}\zeta^2, \tag{A.2}$$

is included as a special case.

Household h has CARA preferences over climate-adjusted consumption,

$$U^h = \mathbb{E}\Big[u^h\big(C_1^h - \phi_h \Gamma\big)\Big], \qquad u^h(y) = -\exp(-\alpha_h y),$$

with $\alpha_h > 0$ and $\phi_h \geq 0$. Writing

$$Y_h \equiv C_1^h - \phi_h h(\bar{X}, \zeta),$$

and using a second–order Taylor expansion of h around $(\bar{X}, 0)$, we obtain

$$h(\bar{X}, \zeta) = h_\zeta(\bar{X}, 0)\zeta + \frac{1}{2} h_{\zeta\zeta}(\bar{X}, 0)\zeta^2 + o(\zeta^2), \tag{A.3}$$

where h_ζ and $h_{\zeta\zeta}$ denote the first and second partial derivatives with respect to ζ. Substituting (A.3) into Y_h yields

$$Y_h = \tilde{C}_1^h - \phi_h h_\zeta(\bar{X}, 0)\zeta - \frac{1}{2}\phi_h h_{\zeta\zeta}(\bar{X}, 0)\zeta^2 + o(\zeta^2),$$

where $\tilde{C}_1^h$ collects terms independent of ζ. To isolate the stochastic components, define the centered square

$$\tilde{\zeta}^2 \equiv \zeta^2 - \mathbb{E}[\zeta^2],$$

and rewrite Y_h as

$$Y_h = \bar{Y}_h + \beta_h^{(1)}\zeta + \beta_h^{(2)}\tilde{\zeta}^2 + o(\zeta^2), \tag{A.4}$$

for suitable constants $\bar{Y}_h, \beta_h^{(1)}, \beta_h^{(2)}$ that depend on $(\alpha_h, \phi_h, \bar{X})$ and on the moments of ζ. Under the assumption that ζ is conditionally normal and that higher-order terms are negligible, the pair $(\zeta, \tilde{\zeta}^2)$ can be treated as a two-dimensional factor vector with mean zero and finite second moments.

The CARA-normal structure implies that each household's portfolio problem again admits a mean-variance representation in terms of Y_h, now with respect to the joint distribution of $(\zeta, \tilde{\zeta}^2)$. Aggregating across households and solving for asset prices leads to an affine SDF of the form

$$m = \exp\big(-A - B^{(1)}\zeta - B^{(2)}\tilde{\zeta}^2\big), \tag{A.5}$$

with factor loadings $B^{(1)}$ and $B^{(2)}$ determined by the distribution of (α_h, ϕ_h) and by the coefficients $h_\zeta(\bar{X}, 0)$ and $h_{\zeta\zeta}(\bar{X}, 0)$.

Suppose now that firm i's dividend takes the reduced-form specification

$$D_i = \mu_i + g_1(Z_i)\zeta + g_2(Z_i)\tilde{\zeta}^2, \tag{A.6}$$

where Z_i collects firm characteristics and g_1, g_2 are measurable functions. A log-linear approximation to the Euler equation $\mathbb{E}[mR_i] = 1$ then yields

$$\mathbb{E}[R_i] - R_f \approx \lambda^{(1)}\frac{g_1(Z_i)}{V_{i,t}} + \lambda^{(2)}\frac{g_2(Z_i)}{V_{i,t}}, \tag{A.7}$$

for factor prices of climate risk $\lambda^{(1)}, \lambda^{(2)}$ that are linear functions of $(B^{(1)}, B^{(2)})$ and of the conditional second moments of $(\zeta, \tilde{\zeta}^2)$. Expression (A.7) shows that allowing for quadratic (or more

general smooth) damages in ζ does not alter the linear factor structure: it simply enlarges the set of priced characteristics from a single normalized function $g(Z_i)/V_i$ to a vector $(g_1(Z_i)/V_i, g_2(Z_i)/V_i)$, corresponding to exposure to the linear and quadratic components of the climate factor. In particular, if the technological channel through which ζ enters cash flows remains proportional to emission levels, as in

$$g_1(Z_i) = \eta_1 X_i, \qquad g_2(Z_i) = \eta_2 X_i,$$

then (A.7) implies that expected excess returns are proportional to X_i/V_i with a combined price of risk $\lambda^{(1)}\eta_1 + \lambda^{(2)}\eta_2$, and the qualitative conclusion that emission levels determine the underlying dollar exposure, so that value-normalized exposure is proportional to X_i/V_i. More generally, the quadratic damage term induces an additional priced characteristic $g_2(Z_i)/V_i$ whose functional form is determined by the way ζ^2 enters the technology and damage processes, but there is no mechanism that generically restores emission intensity X_i/V_i as a unique sufficient statistic.

A.2 Multiple climate shocks: physical, policy, and technology factors

I next allow for multiple climate-related shocks. Let $f_{t+1} = (f_{1,t+1}, \ldots, f_{K,t+1})^\top$ be a vector of K mean-zero factors that may represent, for example, a physical damage factor, a policy or carbon-price factor, and a technology factor. As in Section 4, suppose that firm i's dividend can be written as

$$D_{i,t+1} = \mu_{i,t+1} + \sum_{k=1}^{K} g_k(Z_{i,t}) f_{k,t+1}, \tag{A.8}$$

where $Z_{i,t}$ includes $(X_{i,t}, V_{i,t})$ and any other relevant characteristics, and $g_k(\cdot)$ are exposure functions determined by technology, market structure, and the incidence of policy. Climate damages are now given by a general function

$$\Gamma_t = H(\bar{X}_{t-1}, f_t), \tag{A.9}$$

where H is twice continuously differentiable in its arguments. Under CARA preferences $u^h(y) = -\exp(-\alpha_h y)$, household h's climate-adjusted consumption takes the form

$$Y_{h,t+1} = \tilde{C}_{h,t+1} - \phi_h H(\bar{X}_t, f_{t+1}).$$

A second-order Taylor expansion of H around a baseline $(\bar{X}_t, f_{t+1}) = (\bar{X}_t, 0)$ yields

$$H(\bar{X}_t, f_{t+1}) = H(\bar{X}_t, 0) + \sum_{k=1}^{K} H_k(\bar{X}_t, 0) f_{k,t+1} + \frac{1}{2} \sum_{k,\ell=1}^{K} H_{k\ell}(\bar{X}_t, 0) f_{k,t+1} f_{\ell,t+1} + o(\|f_{t+1}\|^2), \quad \text{(A.10)}$$

where H_k and $H_{k\ell}$ denote the first and second partial derivatives of H with respect to its factor arguments.

Under the maintained small-shock assumption, the leading-order contribution of the quadratic terms in (A.10) to the mean-variance approximation can be summarized by a finite set of additional factors constructed from products $f_{k,t+1} f_{\ell,t+1}$, for example their centered versions $f_{k,t+1} f_{\ell,t+1} - \mathbb{E}_t[f_{k,t+1} f_{\ell,t+1}]$. Aggregating across households and solving the mean-variance portfolio problem then yields an affine SDF of the form

$$m_{t+1} = \exp\Big(- A_t - \sum_{k=1}^{K} B_{k,t} f_{k,t+1} - \sum_{k,\ell=1}^{K} C_{k\ell,t}\, \tilde{f}_{k\ell,t+1} \Big), \quad \text{(A.11)}$$

where $\tilde{f}_{k\ell,t+1}$ are centered quadratic factor terms and the coefficients $B_{k,t}, C_{k\ell,t}$ are functions of (α_h, ϕ_h) and of the derivatives $H_k(\bar{X}_t, 0)$ and $H_{k\ell}(\bar{X}_t, 0)$. Although dividends in (A.8) are linear in the primitive factors $f_{k,t+1}$, the stochastic discount factor in (A.11) depends on both the linear factors and their centered quadratic combinations $\tilde{f}_{k\ell,t+1}$. As a result, the covariance $\mathrm{Cov}_t(m_{t+1}, D_{i,t+1})$ that determines expected returns involves cross-moments between $f_{k,t+1}$ and $\tilde{f}_{k\ell,t+1}$. These interactions generate induced exposure coefficients associated with the quadratic factor terms even though the primitive dividend process is linear. The resulting functions $\tilde{g}_{k\ell}(Z_{i,t})$ therefore represent effective cash-flow exposures arising from the joint dependence of the SDF and dividends on the underlying climate factors, rather than primitive quadratic loadings in dividends themselves. Substituting (A.8) and (A.11) into the Euler equation and applying the same linearization arguments as in Proposition 4 yields

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} \approx \sum_{k=1}^{K} \lambda_{k,t} \frac{g_k(Z_{i,t})}{V_{i,t}} + \sum_{k,\ell=1}^{K} \lambda_{k\ell,t} \frac{\tilde{g}_{k\ell}(Z_{i,t})}{V_{i,t}}, \quad \text{(A.12)}$$

for suitable factor prices of risk $\lambda_{k,t}, \lambda_{k\ell,t}$ and exposure functions $\tilde{g}_{k\ell}(Z_{i,t})$ associated with the quadratic factor terms. Under conditional normality of the primitive factors, the centered quadratic terms are orthogonal to the linear factors and their contribution vanishes, reducing (A.12) to the linear multi-factor pricing equation used in the main text. Expression (A.12) is the natural

multi-factor generalization of (4.5): each physical, policy, or technology factor is associated with a priced characteristic given by its exposure function in the dividend and damage processes, while the quadratic combinations reflect higher-order interactions induced by the dependence of the stochastic discount factor on nonlinear climate damages.

The key observation is that neither the presence of nonlinear damages nor the presence of multiple factors alters the basic structure of our earlier results: to leading order, expected excess returns are linear in a vector of normalized exposure coefficients $(g_1(Z_{i,t})/V_{i,t}, \ldots, g_K(Z_{i,t})/V_{i,t}, \tilde{g}_{k\ell}(Z_{i,t})/V_{i,t})$, where the underlying dollar exposure functions $g_k(Z_{i,t})$ are model-specific functions of firm characteristics $Z_{i,t}$. In particular, if policy shocks enter cash flows through a term of the form $-\tau_{t+1}X_{i,t}$, the corresponding exposure function is naturally $g_k(Z_{i,t}) = -X_{i,t}$, even in the presence of additional physical or technology factors. Conversely, if physical shocks enter via damages that are nonlinear in $\bar{X}_t$ and $f_{k,t+1}$, the resulting priced characteristics will typically involve nonlinear functions of emissions and other firm attributes.

Crucially, the sign structure of the exposure function differs across channels: under the transition benchmark in Section 2, a positive shock $\zeta > 0$ simultaneously raises aggregate damages and brown firms' cash flows, reflecting a more accommodating regulatory environment. A physical damage shock has the opposite transmission: a positive realization reduces firm dividends directly through destruction of capital, supply-chain disruption, or productivity loss, so that $g_k(Z_{i,t}) < 0$ for affected firms regardless of emission levels. These two channels therefore cannot be collapsed into a single scalar shock without breaking one of the two sign structures, which is why the multi-factor representation is the appropriate vehicle for analyzing physical risk in its own right. In neither case is there a general reason for the emission intensity $X_{i,t}/V_{i,t}$ to emerge as a unique sufficient statistic for premia. The qualitative conclusions of Sections 3.4 and 4 are therefore robust to quadratic damages and to multiple climate shocks: the priced objects are value-normalized exposure functions $g_k(Z_{i,t})/V_{i,t}$ determined by technology, policy, and preferences, and these need not collapse to a single ratio such as $X_{i,t}/V_{i,t}$.

B CRRA preferences and incomplete markets

This appendix shows that the main pricing structure derived in the CARA-normal benchmark carries over, in local approximation, to an economy with CRRA preferences and incomplete markets. The key difference is that the aggregate price of climate risk is now governed by the sum of local

risk tolerances, $\sum_h \bar{Y}_h/\gamma_h$, rather than by the sum of inverse CARA coefficients. Throughout, it is useful to keep the distinction between a *price* (dollar) equation and a *return* equation: the primitive exposure is level-based in dividends, while expected returns load on the normalized exposure.

B.1 Setup

Time is $t = 0, 1$. Uncertainty is driven by a scalar climate shock ζ with

$$\mathbb{E}[\zeta] = 0, \qquad \mathrm{Var}(\zeta) = \sigma_\zeta^2 \in (0, \infty),$$

and all random variables are jointly normally distributed conditional on date-0 information.

There are N firms indexed by $i = 1, \ldots, N$ and H households indexed by $h = 1, \ldots, H$. Each firm has one share outstanding, so market clearing requires $\sum_{h=1}^{H} \theta_i^h = 1$ for all i. Firm i has a fixed level of emissions $X_i \geq 0$ and a date-1 dividend given by

$$D_i = \mu_i + \eta X_i \zeta, \tag{B.1}$$

where $\mu_i \in \mathbb{R}$ is a deterministic component and $\eta > 0$ is a technology parameter common across firms. Aggregate emissions are

$$\bar{X} \equiv \sum_{i=1}^{N} X_i.$$

Climate damages at date 1 are

$$\Gamma(\zeta) = \kappa \bar{X} \zeta, \qquad \kappa > 0. \tag{B.2}$$

The scalar ζ is therefore interpreted as a shock to the per-unit cost or value of emissions, in the reduced-form sense discussed in the main text.

Household h has date-0 wealth $W_0^h > 0$ and trades the N stocks and a risk-free asset with gross return $R_f > 0$. Let V_i denote the date-0 price of stock i, and let θ_i^h denote the number of shares of stock i held by household h. Let b^h denote the holding of the risk-free asset. The budget constraint at $t = 0$ is

$$\sum_{i=1}^{N} V_i \theta_i^h + b^h = W_0^h. \tag{B.3}$$

Date–1 consumption of household h is

$$C_1^h(\zeta) = b^h R_f + \sum_{i=1}^{N} \theta_i^h D_i(\zeta). \tag{B.4}$$

Household h experiences disutility from climate damages with intensity $\phi_h \geq 0$ and has CRRA preferences over climate-adjusted consumption:

$$U^h = \mathbb{E}\left[u_h(C_1^h - \phi_h \Gamma(\zeta))\right], \qquad u_h(x) = \frac{x^{1-\gamma_h}}{1-\gamma_h}, \quad \gamma_h > 0. \tag{B.5}$$

Preferences are thus heterogeneous in risk aversion γ_h and climate concern ϕ_h. Damages are not tradable, so markets are incomplete: the non-pecuniary climate risk cannot be fully hedged with traded assets.

We focus on equilibria in which the shocks are small relative to the level of climate-adjusted consumption for each household. This allows for a local mean-variance approximation of expected utility around a deterministic reference point.

B.2 Local mean-variance approximation under CRRA

Define climate-adjusted consumption of household h as

$$Y_h(\zeta) \equiv C_1^h(\zeta) - \phi_h \Gamma(\zeta). \tag{B.6}$$

Using (B.1), (B.2), and (B.4), we can write

$$\begin{aligned} Y_h(\zeta) &= b^h R_f + \sum_{i=1}^{N} \theta_i^h \mu_i + \left(\sum_{i=1}^{N} \theta_i^h \eta X_i - \phi_h \kappa \bar{X} \right) \zeta \\ &= A_h + B_h \zeta, \end{aligned} \tag{B.7}$$

where

$$A_h \equiv b^h R_f + \sum_{i=1}^{N} \theta_i^h \mu_i, \qquad B_h \equiv \sum_{i=1}^{N} \theta_i^h \eta X_i - \phi_h \kappa \bar{X}.$$

Let $\bar{Y}_h \equiv \mathbb{E}[Y_h] = A_h$ denote the conditional mean of Y_h. For the local approximation we assume that the stochastic component $B_h \zeta$ is small relative to $\bar{Y}_h$. A second-order Taylor expansion of $u_h(\cdot)$

around $\bar{Y}_h$ yields

$$u_h(Y_h) \approx u_h(\bar{Y}_h) + u_h'(\bar{Y}_h)(Y_h - \bar{Y}_h) + \frac{1}{2}u_h''(\bar{Y}_h)(Y_h - \bar{Y}_h)^2.$$

Taking expectations and choosing the expansion point $\bar{Y}_h \equiv \mathbb{E}[Y_h]$ implies

$$\mathbb{E}[u_h(Y_h)] \approx u_h(\bar{Y}_h) + \frac{1}{2}u_h''(\bar{Y}_h)\text{Var}(Y_h).$$

Since $\text{Var}(Y_h) = B_h^2\sigma_\zeta^2$, expected utility is approximately

$$\mathbb{E}[u_h(Y_h)] \approx u_h(\bar{Y}_h) + \frac{1}{2}u_h''(\bar{Y}_h)B_h^2\sigma_\zeta^2.$$

Define the local absolute risk aversion of household h at $\bar{Y}_h$ by

$$\tilde{\alpha}_h \equiv -\frac{u_h''(\bar{Y}_h)}{u_h'(\bar{Y}_h)}. \tag{B.8}$$

Thus, under CRRA preferences, the local approximation maps the household's problem into a mean-variance objective with a state-dependent (local) coefficient of absolute risk aversion.

For CRRA utility,

$$u_h'(x) = x^{-\gamma_h}, \qquad u_h''(x) = -\gamma_h x^{-\gamma_h - 1},$$

so

$$\tilde{\alpha}_h = \frac{\gamma_h}{\bar{Y}_h}. \tag{B.9}$$

Dividing the approximation for $\mathbb{E}[u_h(Y_h)]$ by $u_h'(\bar{Y}_h) > 0$ yields the equivalent objective

$$\mathbb{E}[Y_h] - \frac{\tilde{\alpha}_h}{2}\text{Var}(Y_h) = A_h - \frac{\tilde{\alpha}_h}{2}B_h^2\sigma_\zeta^2. \tag{B.10}$$

The constant terms are irrelevant for portfolio choice. Using $Y_h = C_1^h - \phi_h\Gamma(\zeta)$ and $\mathbb{E}[\zeta] = 0$, maximizing (B.10) is thus equivalent, in the local approximation, to maximizing

$$\mathbb{E}[C_1^h - \phi_h\Gamma(\zeta)] - \frac{\tilde{\alpha}_h}{2}\text{Var}[C_1^h - \phi_h\Gamma(\zeta)]. \tag{B.11}$$

Expression (B.11) has the same form as the CARA-normal mean-variance objective in the main text, with the constant absolute risk aversion coefficient α_h replaced by the local coefficient $\tilde{\alpha}_h = \gamma_h/\bar{Y}_h$.

B.3 First-order conditions and equilibrium price of climate risk

Under the local mean-variance objective (B.11), household h chooses $\{\theta_i^h\}_{i=1}^N$ to maximize

$$\sum_{i=1}^{N} \theta_i^h(\mu_i - V_i R_f) - \frac{\tilde{\alpha}_h}{2}\sigma_\zeta^2 \delta_h^2,$$

where

$$\delta_h \equiv \sum_{j=1}^{N} \theta_j^h \eta X_j - \phi_h \kappa \bar{X}. \tag{B.12}$$

In deriving the local first-order conditions, I evaluate $\tilde{\alpha}_h$ at the reference allocation and treat it as fixed, as is standard in local mean-variance approximations. The first-order condition with respect to θ_i^h is

$$\mu_i - V_i R_f - \tilde{\alpha}_h \sigma_\zeta^2 \delta_h \eta X_i = 0, \qquad i = 1, \ldots, N. \tag{B.13}$$

Because $\mu_i - V_i R_f$ is common across households, any household with an interior position in asset i must have the same value of $\tilde{\alpha}_h \sigma_\zeta^2 \delta_h$. Hence there exists a scalar λ such that

$$\mu_i - V_i R_f = \lambda \eta X_i, \qquad \delta_h = \frac{\lambda}{\tilde{\alpha}_h \sigma_\zeta^2}. \tag{B.14}$$

Equation (B.14) is the *price* (dollar) risk-premium relation. Dividing both sides by V_i and using $\mathbb{E}[R_i] = \mathbb{E}[D_i]/V_i$ yields the corresponding *return* representation,

$$\mathbb{E}[R_i] - R_f = \lambda \eta \frac{X_i}{V_i}. \tag{B.15}$$

Thus, as in the CARA benchmark, emission *levels* X_i govern the primitive cash-flow exposure in dividends, while expected *returns* load on the exposure X_i/V_i.

To solve for λ, multiply both sides of (B.12) by η and sum over h:

$$\sum_{h=1}^{H} \sum_{j=1}^{N} \theta_j^h \eta^2 X_j - \kappa \bar{X} \eta \sum_{h=1}^{H} \phi_h = \eta \sum_{h=1}^{H} \delta_h. \tag{B.16}$$

Using market clearing $\sum_{h=1}^{H} \theta_j^h = 1$ for each j, the left-hand side becomes

$$\sum_{j=1}^{N} \eta^2 X_j - \kappa \bar{X} \eta \sum_{h=1}^{H} \phi_h = \eta^2 \bar{X} - \kappa \bar{X} \eta \sum_{h=1}^{H} \phi_h.$$

The right-hand side of (B.16), using (B.14), is

$$\eta \sum_{h=1}^{H} \delta_h = \eta \sum_{h=1}^{H} \frac{\lambda}{\tilde{\alpha}_h \sigma_\zeta^2} = \frac{\eta\lambda}{\sigma_\zeta^2} \sum_{h=1}^{H} \frac{1}{\tilde{\alpha}_h}.$$

Equating both sides and dividing by η yields

$$\eta \bar{X} - \kappa \bar{X} \sum_{h=1}^{H} \phi_h = \frac{\lambda}{\sigma_\zeta^2} \sum_{h=1}^{H} \frac{1}{\tilde{\alpha}_h}.$$

Solving for λ gives the CRRA counterpart to Proposition 2.

Proposition A.1. *Under the assumptions of this appendix and the local mean-variance approximation, the equilibrium price of climate risk is*

$$\lambda = \frac{\sigma_\zeta^2 \bar{X}}{\sum_{h=1}^{H} \frac{1}{\tilde{\alpha}_h}} \Big(\eta - \kappa \sum_{h=1}^{H} \phi_h\Big), \qquad \tilde{\alpha}_h = \frac{\gamma_h}{\bar{Y}_h}. \tag{B.17}$$

Accordingly, expected excess returns satisfy the return representation (B.15)*,*

$$\mathbb{E}[R_i] - R_f = \lambda\eta \frac{X_i}{V_i},$$

and the corresponding price (dollar) relation in (B.14) *can be written equivalently as*

$$\mathbb{E}[D_i] - V_i R_f = \lambda\eta X_i.$$

The denominator in (B.17) can be written as

$$\sum_{h=1}^{H} \frac{1}{\tilde{\alpha}_h} = \sum_{h=1}^{H} \frac{\bar{Y}_h}{\gamma_h},$$

which is the sum of local risk tolerances $\bar{Y}_h/\gamma_h$. Thus, relative to the CARA case, the price of climate risk is governed by aggregate risk tolerance rather than by $\sum_h 1/\alpha_h$. The term in parentheses, $(\eta - \kappa \sum_h \phi_h)$, is unchanged: as in the CARA benchmark, greater aggregate climate concern $\sum_h \phi_h$ reduces, and can in principle overturn, the purely financial component of the brown premium.

B.4 Discussion and relation to the CARA benchmark

Linearity arises from the CARA-normal benchmark and from the local approximation used here; however, the identification of value-normalized exposure as the relevant priced object does not rely on this assumption.

Proposition A.1 shows that, in a CRRA-incomplete-markets setting, the local linear *price* relation

$$\mathbb{E}[D_i] - V_i R_f = \lambda \eta X_i$$

continues to hold, with a price of climate risk λ given by (B.17). The structure is identical to the CARA-normal benchmark: the primitive cash-flow exposure is proportional to the emission level X_i, the numerator scales with $\sigma_\zeta^2 \bar{X}$, and the aggregate climate concern $\sum_h \phi_h$ enters linearly in the term $(\eta - \kappa \sum_h \phi_h)$. The only change is that the denominator is the sum of local risk tolerances, $\sum_h \bar{Y}_h/\gamma_h$, rather than the sum of inverse CARA coefficients.

Translating this into a statement about *returns* requires dividing by the asset price: from the price relation one obtains

$$\mathbb{E}[R_i] - R_f = \frac{\mathbb{E}[D_i] - V_i R_f}{V_i} = \lambda \eta \frac{X_i}{V_i}.$$

Under the local approximation, the economic interpretation of the incomplete-markets equilibrium is unchanged. Heterogeneous risk aversion and climate concern, combined with unhedgeable climate damages, generate a nontrivial price of climate risk λ and a one-factor representation in which the primitive exposure is level-based in dividends while return premia load on the normalized exposure. The CARA and CRRA formulations therefore lead to qualitatively identical asset-pricing implications within the class of small-shock, linear-exposure environments considered in the main text.

C Welfare Losses from Market Incompleteness

This appendix provides the formal derivation of the welfare-loss measure. The purpose is to make precise how welfare comparisons are conducted in the presence of incomplete markets and heterogeneous climate concern, and to clarify the sense in which equilibrium carbon premia differ from planner implementation wedges. As emphasized in the main text, the analysis here is not intended to map observed carbon premia into a social cost of carbon, but rather to characterize the allocation-level efficiency cost that arises when climate risk is imperfectly shared.

C.1 Planner welfare benchmark

Let $W(\cdot)$ denote the planner welfare functional introduced in Section 7.3. For any feasible aggregate allocation (C_0, C_1, X_1),

$$W(C_0, C_1, X_1) \;\equiv\; u(C_0) + \beta\big(u(C_1) - d(X_1)\big),$$

where $d(X_1)$ denotes aggregate climate damages. Let $(C_0^{FB}, C_1^{FB}, X_1^{FB})$ denote the first-best allocation that solves the planner's problem under complete markets, and define

$$W^{FB} \;\equiv\; W(C_0^{FB}, C_1^{FB}, X_1^{FB}).$$

C.2 Incomplete-markets allocation

Let $(C_0^{IM}, C_1^{IM}, X_1^{IM})$ denote the aggregate allocation induced by the competitive equilibrium in the incomplete-markets economy described in Section 3, with the same technology, damage function, and shock distribution as in the planner benchmark. The planner evaluation of this second-best allocation is

$$W^{IM} \;\equiv\; W(C_0^{IM}, C_1^{IM}, X_1^{IM}).$$

Because climate damages are unspanned and households differ in climate concern, the allocation $(C_0^{IM}, C_1^{IM}, X_1^{IM})$ generally does not satisfy the planner's first-order conditions, even though asset prices in equilibrium reflect well-defined carbon premia.

C.3 Consumption-equivalent welfare loss

The welfare loss from market incompleteness is defined in consumption-equivalent units. Specifically, let $\Delta c_1 \geq 0$ denote the unique scalar (on the relevant domain) that satisfies

$$u(C_0^{FB}) + \beta\big(u(C_1^{FB} - \Delta c_1) - d(X_1^{FB})\big) \;=\; W^{IM}.$$

The quantity Δc_1 measures the reduction in date-1 consumption that makes the planner indifferent between the first-best allocation and the incomplete-markets allocation.

For reporting purposes, we also express the welfare loss as a fraction of first-best date-1 con-

sumption,

$$\Delta W \equiv \frac{\Delta c_1}{C_1^{FB}}.$$

C.4 Interpretation and decomposition

The welfare loss ΔW captures two distinct channels. First, incomplete markets prevent full insurance against aggregate climate risk, leading to inefficient consumption smoothing relative to the planner benchmark even in the absence of preference heterogeneity. Second, heterogeneity in climate concern implies that the equilibrium stochastic discount factor reflects a wealth-weighted aggregation of private marginal rates of substitution rather than the planner's marginal valuation of climate damages. As a result, equilibrium investment and emission choices do not internalize the social cost of carbon.

Importantly, ΔW is an allocation object and need not be monotone in the magnitude of equilibrium carbon premia. Carbon premia summarize second-best pricing of exposure to climate and policy shocks, whereas ΔW is pinned down by the gap between allocations under complete and incomplete risk sharing evaluated under a common welfare criterion. This distinction clarifies why observed carbon premia should not be interpreted as sufficient statistics for welfare or for optimal climate policy.

D Infinite-horizon production economy

This section embeds the static incomplete-markets structure into a dynamic production economy. The goal is to show formally that, under standard CARA-normal assumptions, the stochastic discount factor (SDF) admits an exponential-affine representation with a climate-shock loading that is directly linked to the cross-section of (α_h, ϕ_h), and that the resulting dynamic pricing relation preserves the emission-based structure derived in the static model.

Technology, capital, and emissions Time is discrete, $t = 0, 1, 2, \ldots$. Each firm $i \in \{1, \ldots, N\}$ chooses capital $K_{i,t}$ and an emission scale $q_{i,t}$ at each date t. Output at t is given by

$$Y_{i,t} = F_i(K_{i,t}, q_{i,t}) - \Psi_i(q_{i,t}), \tag{D.1}$$

where $F_i : \mathbb{R}_+^2 \to \mathbb{R}_+$ is strictly concave and twice continuously differentiable, and $\Psi_i : \mathbb{R}_+ \to \mathbb{R}_+$ is convex.

Emissions of firm i at date t are

$$X_{i,t} = \chi_i q_{i,t}, \qquad \chi_i > 0, \tag{D.2}$$

so that aggregate emissions are $\bar{X}_t = \sum_{i=1}^{N} X_{i,t}$. Capital evolves according to

$$K_{i,t+1} = (1 - \delta_i)K_{i,t} + I_{i,t}, \qquad \delta_i \in (0, 1). \tag{D.3}$$

Climate risk and dividends Let $\{\zeta_t\}_{t\geq 1}$ be an i.i.d. climate or policy shock with $\mathbb{E}[\zeta_{t+1}|\mathcal{F}_t] = 0$ and $\mathrm{Var}[\zeta_{t+1}|\mathcal{F}_t] = \sigma_\zeta^2$. Climate damages at $t+1$ are

$$\Gamma_{t+1} = \kappa \bar{X}_t \zeta_{t+1}, \qquad \kappa > 0. \tag{D.4}$$

Firm i's dividend at date $t+1$, conditional on decisions at time t, is

$$D_{i,t+1} = F_i(K_{i,t+1}, q_{i,t}) - \Psi_i(q_{i,t}) - I_{i,t+1} + \eta X_{i,t}\zeta_{t+1} + \tilde{\varepsilon}_{i,t+1}. \tag{D.5}$$

where $\tilde{\varepsilon}_{i,t+1}$ is idiosyncratic. As in the static model, the primitive cash-flow exposure to climate risk is proportional to the *level* of emissions chosen in the previous period, $X_{i,t}$, consistent with a one-period implementation lag.

Households, incomplete markets, and dynamic mean–variance representation There are H households indexed by h. Household h has preferences

$$U^h = \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t u^h(C_{h,t} - \phi_h \Gamma_t)\right], \qquad u^h(x) = -\exp(-\alpha_h x), \tag{D.6}$$

with heterogeneity in (α_h, ϕ_h). Climate damages Γ_t are not traded, so markets are incomplete.

Define climate-adjusted consumption

$$\tilde{C}_{h,t} = C_{h,t} - \phi_h \Gamma_t.$$

Under CARA utility and joint conditional normality, the household's dynamic problem reduces at each date t to

$$\max \ \mathbb{E}_t[\tilde{C}_{h,t+1}] - \frac{\alpha_h}{2}\mathrm{Var}_t(\tilde{C}_{h,t+1}), \tag{D.7}$$

which is the exact dynamic analogue of the static mean-variance problem.

Aggregate stochastic discount factor and its climate loading Let $\theta_{i,t}^h$ denote household h's holdings of firm i at date t. Using (D.4)-(D.5), the innovation to climate-adjusted consumption can be written as

$$\tilde{C}_{h,t+1} - \mathbb{E}_t[\tilde{C}_{h,t+1}] = \Big(\sum_{i=1}^{N} \theta_{i,t}^h \eta X_{i,t} - \phi_h \kappa \bar{X}_t\Big)\zeta_{t+1} + \tilde{\varepsilon}_{h,t+1}. \tag{D.8}$$

Using the mean-variance representation (D.7), optimal portfolio choice implies that expected excess returns satisfy

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \alpha_h \operatorname{Cov}_t(R_{i,t+1}, \tilde{C}_{h,t+1}).$$

Substituting (D.5) and (D.8) yields

$$\mathbb{E}_t[D_{i,t+1}] - (1 + R_{f,t})V_{i,t} = \alpha_h \sigma_\zeta^2 \delta_{h,t} \eta X_{i,t}, \qquad \delta_{h,t} = \sum_{j=1}^{N} \theta_{j,t}^h \eta X_{j,t} - \phi_h \kappa \bar{X}_t. \tag{D.9}$$

Since the left-hand side does not depend on h, market clearing implies the existence of a scalar λ_t such that

$$\lambda_t = \frac{\sigma_\zeta^2 \bar{X}_t}{\sum_{h=1}^{H} 1/\alpha_h}\Big(\eta - \kappa \sum_{h=1}^{H} \phi_h\Big). \tag{D.10}$$

SDF representation and pricing implications The aggregate SDF between t and $t+1$ admits the exponential-affine form

$$m_{t+1} = M_t \exp(-B_t \zeta_{t+1}), \tag{D.11}$$

with $\mathbb{E}_t[m_{t+1}] = 1/(1 + R_{f,t})$. A factor representation to the Euler equation implies the payoff risk-premium relation

$$\mathbb{E}_t[D_{i,t+1}] - (1 + R_{f,t})V_{i,t} = B_t \sigma_\zeta^2 \eta X_{i,t}. \tag{D.12}$$

Comparing (D.12) with (D.9) and using (D.10) identifies the SDF loading,

$$B_t = \frac{\lambda_t}{\sigma_\zeta^2} = \frac{\bar{X}_t}{\sum_{h=1}^{H} 1/\alpha_h}\Big(\eta - \kappa \sum_{h=1}^{H} \phi_h\Big). \tag{D.13}$$

Dividing (D.12) by firm value yields the expected excess return,

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = B_t \sigma_\zeta^2 \eta \frac{X_{i,t}}{V_{i,t}}. \tag{D.14}$$

The pricing kernel loading B_t summarizes the aggregate price of climate risk. Although B_t is common across firms, equilibrium emissions $X_{i,t}$ and firm values $V_{i,t}$ are jointly determined by firm-specific technologies, capital stocks, and adjustment costs. Consequently, comparisons based solely on contemporaneous emission intensity $X_{i,t}/V_{i,t}$ need not preserve the ranking of firms by equilibrium exposure to climate risk once emission policies are endogenous and firm values reflect continuation payoffs under optimal adjustment.

Dynamic asset pricing relation Combining (D.14) and (D.13) yields

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \lambda_{X,t} \frac{X_{i,t}}{V_{i,t}}, \qquad \lambda_{X,t} \equiv \lambda_t \eta. \tag{D.15}$$

Equation (D.15) implies that, conditional on equilibrium objects at date t, expected returns are proportional to emission intensity $X_{i,t}/V_{i,t}$. In this reduced-form sense, intensity summarizes cross-sectional differences in expected returns within a given equilibrium. However, this does not imply that emission intensity is a primitive or structurally sufficient statistic for climate exposure. Although equation (D.15) can be written in the one-factor form $\mathbb{E}_t[R_{i,t+1} - R_{f,t}] = \lambda_t g(Z_{i,t})/V_{i,t}$ with reduced-form exposure $g(Z_{i,t})$ proportional to $X_{i,t}$, in the dynamic economy $g(Z_{i,t})$ reflects firms' equilibrium emission policies, which respond endogenously to climate shocks through optimal investment and adjustment decisions. These responses depend on firm-specific state variables–capital stocks, adjustment costs, and expectations of future policy–that differ across firms with identical intensity. The mapping from primitives to equilibrium exposure therefore need not be homogeneous of degree one in $(X_{i,t}, V_{i,t})$, and intensity emerges as a reduced-form normalization rather than a primitive sufficient statistic. This distinction underlies Proposition 3, which characterizes the knife-edge conditions under which intensity is structurally sufficient.

Firms' dynamic investment and emission choices Let $V_{i,t}$ denote the ex-dividend market value of firm i at date t. Shareholders value the stream of dividends using the aggregate stochastic discount factor:

$$V_{i,t} = \mathbb{E}_t \left[\sum_{s=t+1}^{\infty} \left(\prod_{u=t+1}^{s} m_u \right) D_{i,s} \right]. \tag{D.16}$$

Capital evolves according to (D.3) where investment $I_{i,t}$ is chosen at date t and paid immediately.

Dividends at date t are given by

$$D_{i,t} = F_i(K_{i,t}, q_{i,t}) - \Psi_i(q_{i,t}) - I_{i,t} + \eta X_{i,t-1}\zeta_t + \tilde{\varepsilon}_{i,t}. \tag{D.17}$$

Firm i chooses sequences $\{K_{i,t+1}, q_{i,t}\}_{t\geq 0}$ to maximize $V_{i,0}$, taking prices and the SDF process $\{m_{t+1}\}$ as given. The Bellman equation is therefore

$$V_{i,t} = \max_{K_{i,t+1},\, q_{i,t}} \left\{ D_{i,t} + \mathbb{E}_t[m_{t+1} V_{i,t+1}] \right\}. \tag{D.18}$$

We assume that investment and emission choices made at date t affect production, emissions, and dividends realized at date $t+1$ through the capital law of motion (D.3). This one-period implementation lag ensures that the firm's control variables enter next-period cash flows and therefore the Euler equation determining equilibrium asset prices.

The first-order condition with respect to the emission scale $q_{i,t}$ is

$$\mathbb{E}_t\left[m_{t+1} \frac{\partial D_{i,t+1}}{\partial q_{i,t}} \right] + \mathbb{E}_t\left[m_{t+1} \frac{\partial V_{i,t+1}}{\partial q_{i,t}} \right] = 0. \tag{D.19}$$

Under the maintained assumption that $q_{i,t}$ affects future firm value only through next-period dividends, the second term in (D.19) vanishes. Because capital evolves according to (D.3) and does not depend on $q_{i,t}$, emission choices affect continuation value only through dividends realized at $t+1$. Using (D.5), the first-order condition simplifies to

$$\mathbb{E}_t\left[m_{t+1} \left(\frac{\partial F_i(K_{i,t+1}, q_{i,t})}{\partial q_{i,t}} - \Psi_i'(q_{i,t}) + \eta \chi_i \zeta_{t+1} \right) \right] = 0. \tag{D.20}$$

Using (D.11) and normality of ζ_{t+1},

$$\mathbb{E}_t[m_{t+1}\zeta_{t+1}] = -B_t \sigma_\zeta^2 \, \mathbb{E}_t[m_{t+1}],$$

so that (D.20) can be written as

$$\mathbb{E}_t[m_{t+1}] \left(\frac{\partial F_i(K_{i,t+1}, q_{i,t})}{\partial q_{i,t}} - \Psi_i'(q_{i,t}) - \eta \chi_i B_t \sigma_\zeta^2 \right) = 0. \tag{D.21}$$

Dividing by $\mathbb{E}_t[m_{t+1}]$ and using $B_t = \lambda_t/\sigma_\zeta^2$ yields

$$\frac{\partial F_i(K_{i,t+1}, q_{i,t})}{\partial q_{i,t}} - \Psi_i'(q_{i,t}) = \eta \chi_i \lambda_t. \tag{D.22}$$

Equation (D.22) shows that firms equate marginal private benefits of emissions to a climate-risk pricing wedge governed by the same equilibrium price of risk λ_t that prices returns in (D.15). In contrast to a representative-agent benchmark, this wedge reflects equilibrium asset pricing under incomplete markets rather than a planner's Pigouvian tax based on emission intensity (violating Proposition 3 (ii)–(iii)).

Although expected returns admit the reduced-form representation $\mathbb{E}_t[R_{i,t+1} - R_{f,t}] = \eta B_t X_{i,t}/V_{i,t}$, this expression should be interpreted as a pricing relation conditional on equilibrium emission choices. In the dynamic economy, both emissions $X_{i,t}$ and firm values $V_{i,t}$ are endogenous objects determined by optimal investment and adjustment decisions that depend on firm-specific state variables. As a result, the scale-invariance condition required for emission intensity to serve as a robust sufficient statistic becomes dynamically fragile: intensity-based comparisons need not be stable once equilibrium policies and continuation values are taken into account (Proposition 3).

E Additional Tables

Table A.1: Scope 1 Emissions Growth and Future Stock Returns

VARIABLES	(1)	(2)	(3)	(4)
	Stock returns			
	1 month ahead		12 months ahead	
Scope 1 Emissions Growth	0.491***	0.442***	-0.031	-0.056
	(0.057)	(0.054)	(0.034)	(0.034)
Log(Scope 1 Emissions)		0.090***		0.044***
		(0.016)		(0.010)
Log(Size)	-0.212***	-0.259***	-0.140***	-0.163***
	(0.050)	(0.051)	(0.037)	(0.037)
Log(M/B)	-0.258***	-0.218***	-0.157*	-0.138*
	(0.071)	(0.068)	(0.074)	(0.071)
Book leverage	-0.061	-0.166	-0.112	-0.159
	(0.246)	(0.247)	(0.176)	(0.174)
Momentum	0.437**	0.429**	-0.392*	-0.393*
	(0.162)	(0.163)	(0.188)	(0.188)
Investment/Assets	-0.628	-0.497	0.136	0.190
	(0.841)	(0.858)	(0.738)	(0.750)
Log(PPE)	0.088***	0.059**	0.056**	0.042**
	(0.028)	(0.027)	(0.020)	(0.020)
Volatility	1.931	1.932	0.418	0.408
	(3.161)	(3.155)	(1.516)	(1.515)
ROE	0.911***	0.849***	0.546	0.515
	(0.257)	(0.264)	(0.313)	(0.312)
Age	1.128	0.905	-0.436	-0.540
	(0.925)	(0.946)	(0.927)	(0.968)
Constant	2.091***	1.834***	1.800***	1.676***
	(0.449)	(0.464)	(0.297)	(0.294)
Industry×time fixed effects	Yes	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes	Yes
Observations	1,006,227	1,005,779	862,397	862,073
R-squared	0.232	0.232	0.247	0.247

Robust standard errors in parentheses

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$